Innovation Networks and R&D Allocation*

Ernest Liu† *Princeton and NBER*

Song Ma‡ *Yale and NBER*

July 5, 2023

Abstract

We study the cross-sector allocation of R&D resources in a multisector growth model with an innovation network, where one sector's past innovations may benefit other sectors' future innovations. Theoretically, we solve for the optimal allocation of R&D resources. We show a planner valuing long-term growth should allocate more R&D toward central sectors in the innovation network, but the incentive is muted in open economies that benefit more from foreign knowledge spillovers. We derive sufficient statistics for evaluating the welfare gains from improving R&D allocation. Empirically, we build the global innovation network based on patent citations and establish its empirical relevance for knowledge spillovers. We evaluate R&D allocative efficiency across countries using model-based sufficient statistics. Japan has the highest allocative efficiency among the advanced economies. For the U.S., improving R&D allocative efficiency could generate more than 8% welfare gains.

Keywords: Innovation Network, Resource Allocation, Growth
JEL Classification: O33, O38, F43

*Previously circulated as "Innovation Networks and Innovation Policy." We thank Philippe Aghion, Manuel Amador, Pol Antras, David Atkin, Salome Baslandze, Ariel Burstein, Lorenzo Caliendo, Vasco Carvalho, Ben Golub, Jill Grennan, Matt Jackson, Ben Jones, Chad Jones, Hugo Hopenhayn, Bill Kerr, Pete Klenow, Sam Kortum, Atif Mian, Ezra Oberfield, Alessandra Peter, Stephen Redding, Peter Schott, Kjetil Storesletten, Alireza Tahbaz-Salehi, Aleh Tsyvinski, John Van Reenen, Jaume Ventura, Heidi Williams, and Kei-Mu Yi for useful feedback. We also thank seminar participants at Cambridge, EIEF, EUI, Minnesota, NBER ITI, Princeton, UNC Chapel Hill, Yale, ZJU, EAGLS, College de France, SED, Global Economic Networks conference, NBER Summer Institute (Innovation), Richmond Fed, Philadelphia Fed, Stanford, SMU, PBCSF, FOM Conference, UCL, Atlanta Fed, UC Berkeley, MIT, and NSF network conference. Xugan Chen, Tianyu Fan, and Daojing Zhai provided excellent research assistance.

†Department of Economics, Princeton University, and NBER, ernestliu@princeton.edu.

‡Yale School of Management, Yale University, and NBER, song.ma@yale.edu.

1 Introduction

How to foster innovation has long been a central question for economists and policy makers. The discussion has concentrated on the amount of resources invested in research and development (R&D) and the cost of under- or over-investment. But how should these R&D resources be allocated across economic sectors or technological fields? This question is important, policy relevant, yet understudied, and it is the focus of this paper. We ask: How should innovation resources be optimally allocated across sectors to take advantage of cross-sector knowledge spillovers and achieve long-term growth? For example, how many resources should an economy devote to R&D in semiconductors relative to consumer electronics, or chemistry relative to pharmaceutics? How should the optimal R&D allocations differ across countries? How are R&D resources allocated across sectors in the real world, and how much gain does it create to improve cross-sector R&D allocative efficiency?

We answer these questions both theoretically and empirically. The key novelty of our theoretical approach is that we introduce a network perspective into modeling the dynamic spillover structure of innovation. This network captures the notion that one sector's innovation activities require researchers and scientists to build on prior discoveries and knowledge, often from outside their own fields or sectors—a key feature in the innovation process. We solve for the optimal cross-sector allocation of R&D resources and derive model-implied sufficient statistics that can assess the allocative efficiency of R&D in real-world economies. The model is applied to data on more than 30 million global patents from all major economies to assess innovation resource (mis)allocation and potential welfare gains from improving allocative efficiency.

This research has two important motivations. First, cross-sector R&D allocation is an important aspect of many R&D policies, ranging from industrial policies that aim to identify and stimulate a certain set of innovative sectors, to science policy seeking to advance science and harvest long-term value. Second, the cross-sector innovation spillover structure presents a unique opportunity to combine classic endogenous growth theory, recent advances in network methods, and detailed global patent data to answer the research question.

We embed an innovation network into an otherwise canonical multisector endogenous growth model. A finite amount of R&D resources (i.e., scientists) may be deployed across sectors to innovate and improve product quality. One sector's past innovations may subsequently, over a long time path, benefit other sectors' future innovation activities by helping scientists in those sectors innovate more productively. We define the innovation network as the weighted directed graph capturing how one sector's innovation activity benefits from another's past innovation. The state variables of the economy are sectoral knowledge stocks, which reflect the accumulation of past innovations in each sector. Through dynamic spillovers across the network, the state variables

form a dynamical system, in which the evolution of the knowledge stock in each sector depends endogenously on the entire history of resource allocation across all sectors of the economy. The key decision of interest is how to efficiently allocate R&D resources across sectors in the network.

We begin by modeling a closed economy. Our baseline setup adopts an analytically convenient functional form—which we relax later—where cross-sector knowledge spillovers form a log-linear dynamical system. Despite the complexity of dynamic network spillovers, we are able to explicitly solve for the optimal path of cross-sector R&D resource allocation and express the closed-form solution in terms of consumer preferences across sectoral products and sectoral importance in stimulating future innovation through the innovation network. This solution is intuitive; it accounts for: (i) the direct effect of R&D on sectoral output, and (ii) indirect network effects on other sectors through R&D spillovers, discounting benefits that occur far in the future. The optimal R&D allocation is also related to the society's discount rate. A society valuing long-term growth (i.e., with a low discount rate) should allocate more resources toward sectors with fundamental technologies that are central in the innovation network, such as semiconductors. These are technologies that can generate widespread and long-lasting knowledge spillovers to many other sectors, directly or indirectly. By contrast, a short-termist society should allocate more R&D resources toward sectors that immediately benefit consumers but may be peripheral in the innovation network.

Formally, the contribution of each sector's R&D to economic growth is captured by the innovation network's eigenvector centrality, which we call "innovation centrality". We show the innovation centrality vector coincides with the growth-maximizing R&D allocation along a balanced growth path. The optimal R&D allocation chosen by a benevolent planner can be written as a weighted average between the innovation centrality vector and the vector representing consumer preferences over different goods. The former represents the planner's incentives to take advantage of knowledge spillovers for future growth, and the latter represents the planner's incentives to expand knowledge in ways that directly benefit the consumer. A patient planner valuing long-term growth would place a higher weight on the former.

The model also allows us to quantitatively evaluate the cross-sector allocation of R&D resources in the data and calculate the potential welfare gains from adopting the optimal R&D allocations, accounting for transitional dynamics. In consumption-equivalent terms, the welfare gains are proportional (in logs) to the inner product between the optimal R&D allocation vector and the log difference between the optimal and the actual R&D allocation vectors. Hence, this inner product—also known as the relative entropy of the actual R&D allocation from the optimal allocation—is a sufficient statistic to evaluate the potential welfare gains from improving of R&D allocation. This sufficient statistic can be calculated using data on sectoral production, the innovation network, and real-world R&D resource allocation, allowing us to quantitatively assess the

R&D allocative efficiency in the data.

Our baseline results are derived under the tractable formulation where the innovation network is exogenous and invariant to R&D allocations or the levels of knowledge stock, so that the cross-sector knowledge spillovers form a log-linear dynamical system. Under this formulation, the optimal R&D allocation is time-invariant and holds globally along the entire transition path; the sufficient statistic for the welfare gains from R&D reallocation is also global and accounts for the entire transition path. This tractable benchmark may appear restrictive at first, but we show that our analysis holds beyond the log-linear formulation and is more general. In a richer environment with an endogenous innovation network—where cross-sector spillovers depend on the levels of knowledge stock and thus past R&D allocations—our welfare sufficient statistic is still valid to first-order around a balanced growth path. That is, even though the optimal R&D allocation is no longer time-invariant and instead depends on the levels of knowledge stock across sectors, our sufficient statistic serves as a first-order local approximation around a balanced growth path for the welfare gains from reallocating R&D resources; the welfare impact arising from the endogenous changes in the network (due to departure from log-linearity) is second-order in nature.

We extend our model to incorporate an important source of cross-country heterogeneity: knowledge spillovers from abroad. As we show, some countries, like the U.S. and Japan, rely more on domestic knowledge spillovers and less on foreign knowledge spillovers, while other economies benefit more from foreign spillovers particularly from the technologically advanced ones. Intuitively, from the perspective of maximizing domestic welfare, an economy receiving relatively more spillovers from abroad—and less domestically—should allocate less domestic R&D into the network-central sectors. Similarly, when domestic R&D matters less for long-run spillovers, a less efficient domestic R&D allocation is also less consequential for welfare. We demonstrate these intuitions formally by deriving, in an economy receiving foreign spillovers, both (1) the unilaterally optimal R&D allocations and (2) the sufficient statistics for the welfare gains of R&D reallocation.

The tractability of our model lends itself to a large number of theoretical extensions. We demonstrate that our baseline model, which features a simple production structure, can tractably incorporate a production network of input-output linkages. We also embed our innovation network formulation into a semi-endogenous growth setting (Jones 1995, Bloom et al. 2020, Jones 2022). We host several other extensions in the Online Appendix.

Our empirical analysis starts by constructing a global innovation network from over 36 million patents and their citations, collected from over 40 major patent authorities around the world. The data, obtained from Google Patents and originally based on the EPO worldwide bibliographic (DOCDB) data, contain patent-level information on innovations that took place in most

economies between 1976 and 2020. We construct the innovation network as a weighted directed graph using sectors (and country-sectors in our open economy analysis) as nodes and citation shares from one node to another as the edge. Innovation centrality is highly heterogeneous across 131 3-digit international technological classes (IPCs). A handful of IPCs—such as medical science, computing, and semiconductors—are among the most central in the innovation network. Countries vary widely regarding reliance on foreign spillovers: 70% of citations made by U.S. patents are toward other U.S. patents, but most other economies—including China, South Korea, and Germany—are foreign-reliant, with domestic citation shares well below 50%. Within each country, the innovation network only weakly correlates with the input-output production network, such that there is substantial independent variation in both network structures.

To provide evidence that a sector's innovation activities benefit from past innovation in upstream sectors linked through the patent citation network, we extend Acemoglu, Akcigit, and Kerr (2016)—which analyze the U.S. domestic innovation network—using instrumental variables (IVs) and to the global setting. The IVs for past innovation are constructed based on time-varying sectoral exposure to tax-induced user cost of R&D (Wilson, 2009 and Thomson, 2017); they isolate comovements in patent output driven by knowledge spillovers and not by common shocks to connected sectors (Manski, 1993 and Bloom, Schankerman, and Van Reenen, 2013). We find evidence for directional knowledge spillovers: each sector's innovation output responds only to past upstream innovations and does not respond to past innovation from downstream sectors even though they are also connected. We also show that relative to input-output linkages, the innovation network is a significantly stronger channel through which knowledge spillovers occur.

Our main empirical application connects the model-implied optimal R&D allocation with cross-sector R&D allocation in the real world. For each country and time period, we calculate the unilaterally optimal cross-sector R&D allocation. In the U.S., sectors highly ranked in the optimal allocation are primarily those central in the innovation network, such as medical science, semiconductors, and computing devices. We find that the unilaterally optimal allocation differs significantly across countries. For instance, relative to the U.S., Germany and Japan should optimally allocate more R&D resources to vehicle-related innovation, whereas South Korea should invest more R&D in electric communication technique.

The model-implied optimal R&D allocation strongly predicts actual resource allocation in real-world economies, suggesting that our model provides a reasonable way to understand cross-sector R&D allocation. Specifically, we compare the unilaterally optimal R&D allocation against the actual R&D allocation captured using both sectoral R&D expenditure shares and patent output shares in the data. We find that, across many countries—especially for the most innovative ones—sectors that should have more R&D resources do receive more resources.

Nevertheless, the residual misalignment between the optimal and actual allocations trans-

lates into large potential welfare gains from the reallocation of R&D resources to each country's own efficient benchmark. We find Japan has the most efficient R&D allocation. From the R&D allocation in 2010, adopting the optimal allocation could lead to welfare improvements equivalent to raising consumption by 5.64% at every point in time for Japan and 8.04% in the U.S., the two economies that rely the most on domestic knowledge spillovers. Improving R&D allocation could lead to welfare gain of 5.60% in China, 4.24% in South Korea, and 4.09% in Germany. These economies' R&D allocations are less efficient than Japan's, but their domestic R&D is less welfare consequential because they benefit more from foreign spillovers. These cross-country differences in R&D allocative efficiency are qualitatively stable since the 2000s.

It is important to note that a more allocatively efficient economy is not necessarily more innovative in absolute terms. Instead, our notion of cross-sector allocative efficiency reflects the distance from an economy's actual R&D allocation in the data to this economy's own first-best, efficient benchmark. Also note that, by comparing the R&D allocations in the data to the first-best, our notion of allocative efficiency does not require that we take a stance on firms' equilibrium conduct; instead, it directly calculates the welfare impact of reallocating R&D based on the economic environment. Finally, it is worth emphasizing again that our notion of allocative efficiency concerns the relative allocation of R&D resources across sectors and not the aggregate level of R&D.

In the final part of the paper, we discuss over- and under-allocated technology classes in the U.S. Even though providing a full policy recommendation is beyond the paper's scope, the empirical patterns are nevertheless illuminating and show our model's potential to analyze and address more detailed R&D policy issues. For example, our calculation shows that the technology class most relevant to semiconductor technologies (H01) is under-allocated in the U.S. , providing support to the recent U.S. initiatives to accelerate and catalyze the domestic semiconductor sector. We also find under-allocation of R&D in technology classes related to "green innovation" such as waste and pollution management and alternative energy.

This study relates to several strands of existing work. First, our study contributes to a long line of research on knowledge spillovers and innovation policy (Aghion et al., 2005, Bloom et al., 2013, Lucking et al., 2018, Bloom et al., 2019, Jones and Summers, forthcoming, Hopenhayn and Squintani, 2021), particularly in the context of endogenous economic growth (Jones and Williams, 1998, Acemoglu et al., 2018, Akcigit and Kerr, 2018, Atkeson and Burstein, 2019, Garcia-Macia et al., 2019, Bloom et al., 2020, Akcigit et al., 2021, Cai and Tian, 2021, Koenig et al., forthcoming, Akcigit et al., 2022). We contribute to this literature by tackling a key open question: how to optimally allocate R&D resources across sectors in the presence of an innovation network with cross-sector knowledge spillovers.

Relatedly, our study connects to the literature considering cross-sector knowledge linkages,

including Acemoglu et al. (2016), Cai and Li (2019), and Huang and Zenou (2020), and, in an open economy setting, Cai et al. (2022) and Guillard et al. (2021); and cross-country knowledge diffusion (Caballero and Jaffe, 1993, Jaffe et al., 1993, Eaton and Kortum, 1999, 2006, Coe and Helpman, 1995, Coe et al., 2009, Santacreu, 2015, Buera and Oberfield, 2020); see Keller (2004) and Melitz and Redding (2021) for surveys. We contribute to this literature in three ways: first, we build a new endogenous growth model explicitly considering the dynamic and cross-sector spillovers of knowledge; second, our tractable formulation enables us to derive the social optimal R&D allocation and provide simple sufficient statistics for the welfare gains of reallocating R&D optimally, and we introduce foreign-dependence as an important heterogeneity in our open economy setting; third, we construct a global innovation network using patent data around the world, which allows us to empirically study the R&D allocative efficiency in real-world economies.

We also contribute to the fast-growing literature that models network interactions in a general equilibrium setting (Carvalho, 2010, Gabaix, 2011, Acemoglu et al., 2012, Jones, 2011, 2013, Grassi, 2017, Acemoglu et al., 2015, Baqaee, 2018, Lim, 2018, Oberfield, 2018, Liu, 2019, Baqaee and Farhi, 2019, 2020, Chaney, 2018, Taschereau-Dumouchel, 2020, Kleinman et al., 2022, vom Lehn and Winberry, 2022). Particularly related are recent papers that introduce methodologies for dynamic network analysis (Liu and Tsyvinski, 2022, Kleinman et al., forthcoming) and studies on policy interventions targeting specific sectors in static production and strategic networks (Liu, 2019, Galeotti et al., 2020). Relative to this literature, our contribution is to embed an innovation network into a dynamic growth model and study the optimal allocation of R&D resources.

Finally, we contribute to the large literature of resource (mis)allocation (Restuccia and Rogerson, 2008, Hsieh and Klenow, 2009, Jones, 2013, David and Venkateswaran, 2019, Hsieh et al., 2019, Liu, 2019, Baqaee and Farhi, 2020). While this literature focuses on the static misallocation of production resources—potentially due to market distortions, such as taxes, markups and financial frictions—and primarily within-sector across firms, we study the cross-sector allocation of innovation resources, so our analysis is inherently dynamic in nature. Also note that our allocative efficiency measure does not take a stand on firms' equilibrium conduct and, instead, directly calculates the welfare impact of reallocating R&D based on the economic environment.

The rest of the paper is structured as follows. Section 2 has the model and the theoretical results. Section 3 introduces our data. Section 4 describes the global innovation network and provides evidence of its relevance for knowledge spillovers. Section 5 hosts our main empirical application, where we use the model to evaluate cross-sector R&D allocations across countries and time. Section 6 concludes. A separate Online Appendix contains the derivations of the results in the paper, theoretical extensions, and supplementary materials on data and empirical results.

2 Theory: Endogenous Growth with An Innovation Network

We study the implication of R&D resource allocation in a multisector endogenous growth model with an innovation network. We set up the baseline model in Section 2.1 and analyze the optimal allocation of R&D resources across sectors in Section 2.2. Section 2.3 shows how R&D allocation affects long-run growth along the balance growth path. Section 2.4 shows how reallocating R&D resources affects welfare, taking into account the transitional dynamics.

Our baseline model adopts a tractable log-linear formulation featuring an exogenous innovation network. Section 2.5 shows that our results are more general: in a richer environment with flexible functional forms and an endogenous innovation network, our welfare sufficient statistic is still valid to first-order around a balanced growth path.

Section 2.6 extends the baseline model with knowledge spillovers from abroad. In this setting, we derive the unilaterally optimal R&D allocation and the welfare impact of R&D reallocation.

The tractability of our model lends itself to a large number of theoretical extensions. We discuss some of these extensions in Section 2.7, such as incorporating a production network into the model (Section 2.7.1) and how to embed our innovation network formulation into a semi-endogenous growth setting (Section 2.7.2). Section 2.7.3 discuss potential inefficiencies in a stylized decentralized setting. Section 2.7.4 briefly describes other extensions in the Online Appendix.

2.1 Economic Environment of the Baseline Model

Preferences and Production Technology There is a representative consumer with log flow utility and exponential discounting at rate ρ:

$$V_t = \int_t^\infty e^{-\rho(s-t)} \ln y_s \, \mathrm{d}s. \tag{1}$$

The consumption good at each time t is a Cobb-Douglas aggregator over sectoral goods y_{it}, $i = 1, ..., K$:

$$y_t = \prod_{i=1}^K y_{it}^{\beta_i}, \quad \sum_{i=1}^K \beta_i = 1. \tag{2}$$

We refer to β_i as the consumption elasticity of sector i.

Each sectoral good i is a produced linearly from production workers:

$$y_{it} = q_{it}^{\psi} \ell_{it}. \tag{3}$$

We interpret q_{it}^{ψ} as the sectoral productivity. We refer to q_{it} as sector i's knowledge stock at time t. The collection of sectoral knowledge stocks $\{q_{it}\}_{i=1}^K$ are the state variables of the economy. The exponent ψ parametrizes how sectoral knowledge translates into sectoral productivity.

The Innovation Process R&D expands the knowledge stock.

At time t, mass s_i of R&D resources (e.g., scientists) employed in sector i generate a flow of new innovation n_{it}:

$$n_{it} = s_{it}\eta_i\chi_{it}, \qquad \chi_{it} \equiv \prod_{j=1}^{K} q_{jt}^{\omega_{ij}}, \qquad \sum_{j=1}^{K}\omega_{ij} \equiv 1. \tag{4}$$

η_i is the exogenous component of innovation productivity, and χ_{it} is the endogenous component. Importantly, χ_{it} is a Cobb-Douglas combination of knowledge stock across all sectors. The aggregator χ_{it} implies that a larger knowledge stock q_j in sector j facilitates innovation production in sector i with elasticity ω_{ij}, thereby making scientists in sector i conduct R&D more productively. Our formulation thus captures cross-sector knowledge spillovers; that is, scientists stand on the shoulders of giants across all sectors of the economy. We impose the assumption that χ_{it} has constant returns to scale ($\sum_{j=1}^{K}\omega_{ij} = 1$), which implies sustained and nonexplosive growth. Absent cross-sector knowledge spillovers, $\omega_{ij} = 1$ if $i = j$ and is zero otherwise.

New innovation n_{it} expands the knowledge stock according to the following law of motion:

$$\dot{q}_{it}/q_{it} = \lambda \ln\left(n_{it}/q_{it}\right). \tag{5}$$

The key distinction between n_{it} and q_{it} is that the former is a *flow* variable reflecting R&D activities at time t, whereas the latter is a *stock* variable reflecting the accumulation of past innovations. The rate at which knowledge expands in sector i is increasing in the flow of new innovation and decreasing in the existing knowledge stock q_{it}, capturing the notion that innovation gets harder as the knowledge stock in sector i expands. λ parametrizes the sensitivity of knowledge growth to the flow of new innovation relative to the existing stock.

Throughout the rest of the paper, we use boldface variables to denote column vectors (lowercase) and matrices (uppercase). Let $\boldsymbol{q}_t$ denote the column vector whose i-th entry is q_{it}; $\boldsymbol{q}_t$ captures the economy's state variables.

Definition 1. (**Innovation Network**) The innovation network $\boldsymbol{\Omega} \equiv [\omega_{ij}]$ is the $K \times K$ matrix whose ij-th entry is ω_{ij}.

A key object of this study, the $\boldsymbol{\Omega}$ matrix represents a weighted directed graph in which economic sectors are the graph nodes. Elements of the $\boldsymbol{\Omega}$ matrix ω_{ij} capture the elasticity to which sector i's innovation production benefits from sector j's existing knowledge stock. We refer to sector j as *upstream* to sector i and, conversely, i as *downstream* to j; this terminology captures the notion that knowledge flows from upstream sector j to downstream sector i. Absent cross-sector knowledge spillovers, $\boldsymbol{\Omega} = \boldsymbol{I}$ is the identity matrix. The construction is not limited by any specific sector definition; for instance, innovation networks can be constructed across industrial sectors, technology classes, and scientific fields.

Resources We close the model with resource constraints. The economy is endowed with two exogenous stocks of resources: production workers of mass $\bar{\ell}$ and research scientists of mass $\bar{s}$. Workers are employed to produce sectoral goods as in (3). Scientists are employed to conduct R&D. The market clearing conditions for production workers and scientists are:

$$\sum_{i=1}^{K} \ell_{it} = \bar{\ell}, \qquad \sum_{i=1}^{K} s_i = \bar{s}. \tag{6}$$

Remark. In the baseline model, we separate R&D and production resources for expositional simplicity. As we show below, our results concerns the cross-sector allocation shares of R&D resources ($s_{it}/\bar{s}$), and our characterization is invariant to the level of R&D resources $\bar{s}$. Hence, all of our results hold even in a richer model with factor mobility between R&D and production (see Online Appendix B.10).

2.2 Optimal Allocation of R&D Resources

In this section we characterize the optimal allocation of R&D resources in the economy. Consider a benevolent social planner who chooses the entire time path of worker and scientist allocations across sectors to maximize consumer utility. We can write the planner's problem as

$$V^*\left(\{q_{i0}\}\right) \equiv \max_{\{\ell_{it}, s_{it}\}} \int_0^\infty e^{-\rho t} \sum_{i=1}^{K} \beta_i \ln y_{it} \, \mathrm{d}t, \tag{7}$$

subject to the sectoral aggregator (3) for y_{it}, the flow of new innovation (4), law of motion for sectoral knowledge (5), and the resource constraints (6).

First, recognize from equation (3) that the planner's objective is log-linear in the allocation of production workers, implying the following lemma.

Lemma 1. *The planner allocates production workers in proportion to the consumption elasticity vector $\boldsymbol{\beta}$: for all t, $\ell_{it} = \beta_i \bar{\ell}$ for each sector i and variety ν.*

Lemma 1 simplifies the planner's problem into choosing how to allocate scientists only. Recall $\boldsymbol{\Omega} \equiv [\omega_{ij}]$ is the matrix that encodes the innovation network, and $\ln \boldsymbol{q}_t \equiv [\ln q_{it}]_{i=1}^K$ is the vector of log-knowledge stock at time t. Let $\gamma_{it} \equiv s_{it}/\bar{s}$ denote the share of scientists allocated to sector i at time t, and let $\boldsymbol{\gamma}_t$ denote the vector $[\gamma_{it}]_{i=1}^K$ that sums to one. Using equation (3) to express consumption in terms of production worker allocation and then applying Lemma 1, we rewrite the planner's problem in vector form as

$$\max_{\{\boldsymbol{\gamma}_t\} \text{ s.t. } \boldsymbol{\gamma}_t' \mathbf{1} = 1 \forall t} \psi \cdot \int_0^\infty e^{-\rho t} \boldsymbol{\beta}' \ln \boldsymbol{q}_t \, \mathrm{d}t \tag{8}$$

$$\text{s.t. } \mathrm{d} \ln \boldsymbol{q}_t / \mathrm{d}t = \lambda \cdot \left(\ln \boldsymbol{\eta} + \ln \bar{s} + \ln \boldsymbol{\gamma}_t + (\boldsymbol{\Omega} - \boldsymbol{I}) \ln \boldsymbol{q}_t\right). \tag{9}$$

We obtain (9) by substituting the innovation production function (4) into $\boldsymbol{q}_t$'s law of motion (5).

The planner's problem may seem intractable: the economy features a vector of state variables (sectoral knowledge stocks), and the law of motion involves dynamic network spillovers across sectors, meaning the allocation of R&D in any sector at any time affects the evolution of all state variables in all future times. Our formulation, however, is especially tractable: both the planner's objective function (8) and the law of motion (9) are log-linear in the state variables $\boldsymbol{q}_t$. Such tractability enables us to characterize the solution—the entire time path of optimal R&D allocation—in closed form. Later in Section 2.5 we generalize our analysis to a nonlinear setting.

Proposition 1. *Starting from any vector of initial knowledge stock $\boldsymbol{q}_0$, the optimal R&D allocation is time-invariant and follows, along the entire time path,*

$$\boldsymbol{\gamma}' = \frac{\rho}{\rho+\lambda}\boldsymbol{\beta}'\left(\boldsymbol{I} - \frac{\boldsymbol{\Omega}}{1+\rho/\lambda}\right)^{-1}. \tag{10}$$

Proposition 1 shows the optimal cross-sector R&D allocation is time-invariant and follows $\boldsymbol{\gamma}' \propto \boldsymbol{\beta}'\left(\boldsymbol{I} - \frac{\boldsymbol{\Omega}}{1+\rho/\lambda}\right)^{-1}$; the proportionality constant, $\frac{\rho}{\rho+\lambda}$, ensures that $\boldsymbol{\gamma}$ sums to one. To understand the intuition for the result, note that another way to write the optimal allocation vector of R&D resources $\boldsymbol{\gamma}'$ is:

$$\boldsymbol{\gamma}' \;\propto\; \boldsymbol{\beta}'\sum_{m=0}^{\infty}\left(\frac{\boldsymbol{\Omega}}{1+\rho/\lambda}\right)^{m} = \boldsymbol{\beta}'\left(\boldsymbol{I} + \frac{\boldsymbol{\Omega}}{1+\rho/\lambda} + \left(\frac{\boldsymbol{\Omega}}{1+\rho/\lambda}\right)^{2} + \cdots\right).$$

That is, the Leontief inverse $\left(\boldsymbol{I} - \frac{\boldsymbol{\Omega}}{1+\rho/\lambda}\right)^{-1}$ in (10) can be written as a power series of $\frac{\boldsymbol{\Omega}}{1+\rho/\lambda}$. The first term in the infinite summation, $\boldsymbol{\beta}'\boldsymbol{I} = \boldsymbol{\beta}'$, captures how each sector's knowledge stock directly impacts consumer welfare through product quality. This term coincides with the optimal allocation of production workers (Lemma 1). The products between $\boldsymbol{\beta}'$ and subsequent terms in the power series capture the indirect effect of knowledge creation on consumer welfare, through future innovations and product quality improvements in network-connected sectors. Innovations in sector j benefit sector i by endogenously raising the efficiency of subsequent R&D in sector i, captured by the aggregator χ_{it} in equation (4) with elasticity ω_{ij}, which is the ij-th entry of the innovation network matrix $\boldsymbol{\Omega}$. Improved innovation efficiency in sector i further generates additional knock-on effects, as new knowledge in sector i facilitates future innovations in all sectors that benefit from sector i's knowledge stock; the higher-powered terms in the infinite summation capture these indirect effects.

Because network spillovers occur through sectoral knowledge stock, the flow of new knowledge through current R&D activities can only affect innovative efficiency and product quality in the future. Hence, the importance of network effects in the optimal R&D allocation is modulated by the discount rate ρ relative to λ, the sensitivity of knowledge growth to the flow of new innovation relative to the existing stock. The former (ρ) captures discounting of the future, and the latter (λ) captures how quickly those future benefits materialize. We refer to ρ/λ as the society's

effective discount rate, which is a key parameter in determining optimal innovation allocation. When ρ/λ is high, the planner discounts future benefits heavily, and the network effects play a smaller role. In the limit as $\rho/\lambda \to \infty$, the planner becomes myopic, and the optimal R&D allocation is fully pinned down by the consumer preferences $\boldsymbol{\beta}$. Conversely, a more patient (low ρ/λ) planner allocates more R&D resources to sectors that benefit more sectors in the future, directly or indirectly. Proposition 1 implies that a patient planner directs R&D into basic science; an impatient planner directs R&D into consumer goods that may be peripheral in the innovation network, such as textiles and food products.[1]

2.3 R&D Allocation and Economic Growth

In this section we show how R&D allocation affects the economic growth rate along a balanced growth path (BGP). We demonstrate that the network's eigenvector centrality—what we call "innovation centrality"—is a sufficient statistic for evaluating the growth rate along a BGP and coincides with the growth-maximizing R&D allocation. We show the socially optimal R&D allocation $\boldsymbol{\gamma}$ converges to this growth-maximizing allocation when the planner is infinitely patient.

Definition 2. (Innovation Centrality) The vector of sectoral *innovation centrality*, $\boldsymbol{a} \equiv [a_i]_{i=1}^K$, is the dominant left-eigenvector of the innovation network $\boldsymbol{\Omega}$ with an associated eigenvalue of one, satisfying $\boldsymbol{a}' = \boldsymbol{a}'\boldsymbol{\Omega}$ and $\sum_{i=1}^K a_i = 1$.

Because $\boldsymbol{\Omega}$ is a row-stochastic matrix, the innovation centrality vector $\boldsymbol{a}$ exists and is generically unique by the Perron-Frobenius theorem. We now show $\boldsymbol{a}$ is a key determinant of the BGP growth rate and coincides with the growth-maximizing R&D allocation.

Let $\boldsymbol{b}$ denote a generic vector of allocation shares with nonnegative entries and $\sum_{i=1}^K b_i = 1$.

Lemma 2. *Consider a BGP in which the aggregate consumption grows at a constant rate, with time-invariant allocations of production and R&D resources. Suppose R&D resources are allocated according to the vector* $\boldsymbol{b}$ *(i.e.,* $s_i/\bar{s} = b_i$*); then, along the BGP, the growth rate of knowledge stock is the same across sectors and equals to*

$$g^q(\boldsymbol{b}) = const + \lambda \cdot \boldsymbol{a}' \ln \boldsymbol{b}, \tag{11}$$

where the constant on the right-hand side is equal to $\lambda \cdot (\ln \bar{s} + \boldsymbol{a}' \ln \boldsymbol{\eta})$*. The growth rate of theaggregate consumption is*

$$g^y(\boldsymbol{b}) = \psi \cdot g^q(\boldsymbol{b}). \tag{12}$$

Lemma 2 analytically expresses the BGP growth rate of knowledge stock and the aggregate consumption as functions of the R&D allocation, $\boldsymbol{b}$. Along the BGP, the knowledge stock grows

[1] In Section B.1 of the Online Appendix, we provide an example with three sectors, and we analytically express the optimal allocation based on network structure and effective discount rate ρ/λ.

at the same rate $g^q(\boldsymbol{b})$ across sectors. The endogenous component of $g^q(\boldsymbol{b})$ is λ times the inner product between the innovation centrality ($\boldsymbol{a}$) and the vector of log-R&D allocation shares ($\ln \boldsymbol{b}$) (recall λ is the sensitivity of knowledge growth to the flow of new innovation relative to the existing stock). The exogenous component (the constant term) on the right-hand side of (11) shows that the growth rate is higher when R&D resources are more abundant (higher $\bar{s}$), when R&D leads to more new innovation flows (higher $\boldsymbol{\eta}$), and when a higher flow innovation (relative to the existing stock) leads to faster knowledge growth (higher λ). Since ψ parametrizes how knowledge translates into productivity, (c.f. equation 3), the growth rate of the aggregate consumption is simply ψ times the growth rate of knowledge stock across sectors.

Corollary 1. *The R&D allocation that maximizes the BGP consumption growth rate coincides with the innovation centrality $\boldsymbol{a}$, as it solves the following problem: $\boldsymbol{a} = \arg\max_{\boldsymbol{b}} \boldsymbol{a}' \ln \boldsymbol{b}$, s.t. $\boldsymbol{b} \geq 0, \mathbf{1}'\boldsymbol{b} = 1$.*

This corollary highlights that innovation centrality $\boldsymbol{a}$ coincides with the growth-maximizing R&D allocation along a BGP. Intuitively, a_i captures the extent to which sector i's R&D activities contribute to economic growth, taking into account the network effects. Sectors with higher innovation centrality represent more fundamental technologies in the innovation network.

The corollary also demonstrates that the social planner does not necessarily choose the R&D allocation that maximizes growth. Unlike the socially optimal allocation $\boldsymbol{\gamma}$, which depends on the effective discount rate ρ/λ, the growth-maximizing allocation is equal to the innovation centrality and is independent of these parameters. Intuitively, the social planner maximizes the welfare of the consumer, who may prefer better quality products in the near future from consumption-intensive sectors (e.g., consumer goods such as textiles and food products), and knowledge in these sectors may not generate much knowledge spillovers for future innovations.

One can rewrite the optimal R&D allocation vector $\boldsymbol{\gamma}$ as the solution to the following fixed point equation, which demonstrates how $\boldsymbol{\gamma}$ varies with the effective discount rate ρ/λ:

$$\frac{\rho}{\lambda}\left(\boldsymbol{\gamma}' - \boldsymbol{\beta}'\right) + \boldsymbol{\gamma}'\left(\boldsymbol{I} - \boldsymbol{\Omega}\right) = \mathbf{0}'. \tag{13}$$

Equation (13) demonstrates that the optimal R&D allocation $\boldsymbol{\gamma}$ trades off between consumer preferences $\boldsymbol{\beta}$ and long-run growth captured by the innovation centrality $\boldsymbol{a}$, and ρ/λ modulates the relative importance of these two terms. When ρ/λ is large—an impatient planner—consumer preferences dominate ($\lim_{\rho/\lambda\to\infty} \boldsymbol{\gamma} = \boldsymbol{\beta}$). The planner places more weight on growth as ρ/λ declines; in the limit $\rho/\lambda \to 0$, equation (13) implies that $\boldsymbol{\gamma}'\left(\boldsymbol{I} - \boldsymbol{\Omega}\right) \to \mathbf{0}'$, thus the optimal allocation converges to the growth maximizing allocation ($\lim_{\rho/\lambda\to 0} \boldsymbol{\gamma} = \boldsymbol{a}$).

Proposition 2. *As the planner becomes infinitely patient ($\rho/\lambda \to 0$), the optimal R&D allocation converges to the innovation centrality, which is the growth-maximizing allocation: $\lim_{\rho/\lambda\to 0} \boldsymbol{\gamma} = \boldsymbol{a}$.*

As the planner becomes infinitely impatient ($\rho/\lambda \to \infty$), the optimal R&D allocation converges to the consumption elasticity vector: $\lim_{\rho/\lambda \to \infty} \boldsymbol{\gamma} = \boldsymbol{\beta}$.

2.4 The Impact of R&D Allocation on Welfare

We now derive the welfare impact of R&D allocation, taking into account the transition dynamics.

Proposition 3. *For any initial knowledge stock $\boldsymbol{q}_0$ and any path of worker allocation $\{\boldsymbol{\ell}_t\}$, the difference in consumer welfare generated by two time-invariant R&D allocations $\widetilde{\boldsymbol{b}}$ and $\boldsymbol{b}$ is*

$$V\left(\boldsymbol{q}_0; \{\boldsymbol{\ell}_t\}, \widetilde{\boldsymbol{b}}\right) - V\left(\boldsymbol{q}_0; \{\boldsymbol{\ell}_t\}, \boldsymbol{b}\right) = \frac{\psi\lambda}{\rho^2}\boldsymbol{\gamma}'\left(\ln \widetilde{\boldsymbol{b}} - \ln \boldsymbol{b}\right). \tag{14}$$

Proposition 3 shows that the welfare differences resulting from two R&D allocation vectors can be expressed as the inner product between the optimal R&D allocation $\boldsymbol{\gamma}$ and the log-differences in R&D allocation vectors, multiplied by the scalar $\psi\lambda/\rho^2$. The result holds for any time path of worker allocation and any initial knowledge stock $\boldsymbol{q}_0$; hence, the Proposition can be used for welfare evaluation of policy counterfactuals that reallocate R&D resources across sectors.

Definition 3. (Consumption-Equivalent Welfare Gains from Adopting the Optimal R&D) Consider an economy with time-invariant R&D allocation $\boldsymbol{b}$ and the associated consumption path $\{y_t\}_{t\geq 0}$. The *consumption-equivalent welfare gains from adopting the optimal R&D allocation* $\boldsymbol{\gamma}$, is the scalar $\mathcal{L}(\boldsymbol{b})$ such that the consumer is indifferent between the consumption path $\{\mathcal{L}(\boldsymbol{b}) \times y_t\}_{t\geq 0}$ and the consumption path generated by reallocating R&D optimally according to the vector $\boldsymbol{\gamma}$, while holding worker allocation unchanged.

The scalar $\mathcal{L}(\boldsymbol{b})$ quantifies the welfare impact of reallocating R&D resources from $\boldsymbol{b}$ to $\boldsymbol{\gamma}$ across sectors. Note that, because the flow output is log-additive in the knowledge stock $\boldsymbol{q}_t$ and worker allocation $\boldsymbol{\ell}_t$ (equation 3), $\mathcal{L}(\boldsymbol{b})$ does not depend on the path of worker allocation. We next provide a simple formula for $\mathcal{L}(\boldsymbol{b})$.

Proposition 4. *The consumption-equivalent welfare impact of adopting the optimal R&D is*

$$\mathcal{L}(\boldsymbol{b}) = \exp\left(\frac{\psi\lambda}{\rho}\boldsymbol{\gamma}'\left(\ln \boldsymbol{\gamma} - \ln \boldsymbol{b}\right)\right). \tag{15}$$

Proposition 4 shows that the consumption-equivalent welfare impact of reallocating R&D resources optimally, in logs, the inner product between the optimal allocation $\boldsymbol{\gamma}'$ and the log-difference between the optimal and the actual allocations, $(\ln \boldsymbol{\gamma} - \ln \boldsymbol{b})$, multiplied by ψ, the elasticity of productivity to knowledge stock, and λ/ρ, the inverse of the effective discount rate. The inner product term, also known as the relative entropy, is a statistical distance measure of how $\boldsymbol{b}$ differs from $\boldsymbol{\gamma}$. Note that ρ/λ affects welfare through not only the proportionality constant but also the optimal allocation $\boldsymbol{\gamma}$ (as in Proposition 1).

2.5 General Functional Forms and Endogenous Innovation Network

The baseline model features an exogenous innovation network $\mathbf{\Omega}$, as the elasticities of each sector's innovation productivity to another sector's knowledge stock ($\omega_{ij} \equiv \frac{\partial \ln \chi_{it}}{\partial \ln q_{jt}}$) are exogenous structural parameters. The knowledge spillover dynamics (9) in the baseline model thus form a log-linear dynamical system. Such log-linearity—along with the Cobb-Douglas preferences across goods, represented by the exogenous consumption elasticities $\boldsymbol{\beta}$—makes the baseline model especially tractable. Under this formulation, the optimal R&D allocation is time-invariant and holds globally along the entire transition path; the sufficient statistic for the welfare gains from R&D reallocation is also global and accounts for the gains along the entire transition path.

The log-linearity in the baseline model may appear restrictive at first, as it rules out the possibility that degree of knowledge spillovers depend endogenously on the levels of sectoral knowledge stock: as sector j accumulates knowledge q_{jt}, its contribution for sector i's innovation ($\frac{\partial \ln \chi_{it}}{\partial \ln q_{jt}}$) may rise or fall. Moreover, sector j's importance for consumption ($\frac{\partial \ln y_t}{\partial \ln y_{jt}}$) may also change endogenously. Such nonlinearity can be incorporated by modeling the consumption and innovation spillover elasticities ($\boldsymbol{\beta}$ and $\mathbf{\Omega}$) not as structural parameters but as objects that endogenously depend on the levels of knowledge stock. In such a richer, nonlinear environment, the optimal R&D allocation is no longer time-invariant; instead, it varies with the levels of knowledge stock across sectors (and depends also on the exogenous component η_i of each sector's innovation productivity).

We now show that our baseline results are actually more general. In nonlinear economic environments described above, our sufficient statistic serves as a first-order local approximation around a balanced growth path (BGP) for the potential welfare impact of reallocating R&D resources. The welfare impact arising from the endogenous changes in the consumption elasticities or the network structure (due to departure from log-linearity) is second-order in nature.

Specifically, consider replacing the log-linear consumption aggregator $y_t = \prod_{i=1}^{K} y_{it}^{\beta_i}$ in (2) with a general constant-return-to-scale function $y_t = \mathcal{Y}(\{y_{it}\})$, and suppose the law of motion for sectoral knowledge stock follows

$$\dot{q}_{it}/q_{it} = \lambda \ln\left(s_{it}\mathcal{X}_i\left(\{q_{jt}\}\right)\right) \tag{16}$$

where the R&D spillover function $\mathcal{X}_i(\cdot)$ satisfies homogeneous-of-degree-zero with positive cross-sector spillovers ($\partial \ln \mathcal{X}_i(\cdot)/\partial \ln q_{jt} > 0$ for $i \neq j$). We assume the spillovers are bounded, and without loss of generality we set the bound to one ($|\partial \ln \mathcal{X}_i(\cdot)/\partial \ln q_{jt}| \leq 1 \forall i, j$).[2]

In this environment, consumer welfare given the path of R&D allocation $\{b_{it} \equiv s_{it}/\bar{s}\}$ is

$$\int_0^\infty e^{-\rho t} \ln \mathcal{Y}\left(\left\{q_{it}^{\psi}\ell_{it}\right\}\right) \mathrm{d}t, \quad \text{s.t. } \mathrm{d}\ln q_{it}/\mathrm{d}t = \lambda \cdot \left(\ln\left(\eta_i b_{it}\bar{s}\right) + \ln \mathcal{X}_i\left(\{q_{jt}\}\right)\right) \quad \forall i, \tag{17}$$

[2] In Appendix A.7, we further generalize the law of motion by dropping the functional form assumption on R&D allocation and requiring only that the growth rate of q_{it} is log-separable in R&D allocation and spillovers.

where ℓ_{it} is the measure of production workers allocated to each variety in sector i at time t. In this environment, we can define the consumption and innovation spillover elasticities as

$$\beta_{it} \equiv \frac{\partial \ln \mathcal{Y}\left(\{y_{it}\}\right)}{\partial \ln y_{it}}, \qquad \omega_{ijt} \equiv \begin{cases} \frac{\partial \ln \chi_i(\{q_{kt}\})}{\partial \ln q_{jt}} & \text{if } i \neq j, \\ 1 + \frac{\partial \ln \chi_i(\{q_{jt}\})}{\partial \ln q_{it}} & \text{otherwise.} \end{cases}$$

These elasticities are not structural parameters; instead, they are reduced-form objects evaluated at specific levels of allocation and knowledge stock. Precisely because these elasticities change endogenously, the optimal R&D allocation is no longer time-invariant and instead depends on the levels of knowledge stock across sectors.

Our results in the baseline model extends to this nonlinear environment locally around a BGP, where the allocations of worker and R&D resources are time invariant. Specifically, the consumption and innovation spillover elasticities $\boldsymbol{\beta} \equiv [\beta_i]$ and $\boldsymbol{\Omega} \equiv [\omega_{ij}]$ are also time-invariant in a BGP; hence, we can define $\boldsymbol{\gamma}' \propto \boldsymbol{\beta}' \left(I - \frac{\boldsymbol{\Omega}}{1+\rho/\lambda}\right)^{-1}$ and scale its entries so that $\sum_i \gamma_i = 1$. In this nonlinear environment, the vector $\boldsymbol{\gamma}$ no longer represents the globally optimal R&D allocation (as in the baseline model); instead, it is proportional to the local elasticity of how R&D allocation affects welfare, as shown in the following result.

Proposition 5. *Consider an economy in a balanced growth path with R&D allocation $\boldsymbol{b}$. As a first-order approximation around the initial BGP, the consumption-equivalent welfare gain of moving from allocation $\boldsymbol{b}$ to $\widetilde{\boldsymbol{b}}$, accounting for transitional dynamics, is* $\exp\left(\frac{\psi\lambda}{\rho}\boldsymbol{\gamma}'\left(\ln\widetilde{\boldsymbol{b}} - \ln\boldsymbol{b}\right)\right)$.

The Proposition shows that the sufficient statistic for the welfare impact of R&D reallocation continues to hold as a first-order approximation around a BGP of the general nonlinear environment. Formally, the vector $\boldsymbol{\gamma}$ is proportional to the directional (Gateaux) derivative of welfare with respect to the R&D allocation. The welfare impact arising from the endogenous changes in the consumption elasticities or the network structure (due to departure from log-linearity) is second-order in nature.

2.6 Knowledge Spillovers from Abroad

We will later use our model to assess R&D allocations in real-world economies. As we show, some countries, like the U.S. and Japan, rely more on domestic knowledge spillovers and less on foreign knowledge spillovers, while other economies benefit more from foreign spillovers particularly from the technologically advanced ones. We now extend our model to incorporate this important cross-country heterogeneity, namely the reliance on knowledge spillovers from abroad, and examine how it affects the welfare impact of domestic R&D allocation.

2.6.1 Setup with Knowledge Spillovers from Abroad

We extend the closed economy model in Section 2.1 by introducing international knowledge spillovers and trade. The consumer now values both domestic and foreign goods with the constant-returns-to-scale preference aggregator $\mathcal{C}(\cdot)$:

$$V = \int_0^\infty e^{-\rho t} \ln \mathcal{C}\left(c_t^d, c_t^f\right) \mathrm{d}t,$$

where c_t^d is consumption of domestic goods and c_t^f is consumption of foreign goods. As in the closed economy model, domestic goods represent a Cobb-Douglas aggregation over sectoral goods (equations 2 and 3). The economy can import foreign goods c_t^f by exporting unconsumed domestic goods $\left(y_t - c_t^d\right)$. Let p_t^f denote the relative prices of foreign goods. We impose trade balance:

$$p_t^f c_t^f = y_t - c_t^d. \tag{18}$$

Domestic innovation production benefits from foreign knowledge spillovers $\left\{q_{jt}^f\right\}_{j=1}^K$:

$$n_{it} = s_{it}\eta_i\chi_{it}, \qquad \text{where } \chi_{it} = \prod_{j=1}^K \left[(q_{jt})^{x_{ij}} \left(q_{jt}^f\right)^{1-x_{ij}}\right]^{\omega_{ij}}. \tag{19}$$

Like in the closed economy counterpart (4), n_{it} is the flow of new innovation generated by R&D resources s_{it} in sector i at time t; new innovation leads knowledge accumulation according to (5). χ_{it} is again the endogenous component of R&D efficiency in sector i. The difference here is that domestic R&D in sector i benefits from not only domestic knowledge q_{jt} in sector j but also foreign knowledge q_{jt}^f. The exponent x_{ij} captures the share of domestic contribution of knowledge spillover from sector j to sector i; when $x_{ij} = 1$ for all i, j, the innovation production function (19) coincides with the closed economy version (4). x_{ij}'s could differ across countries for various reasons. For instance, it could reflect a country's R&D comparative advantage; it could also reflect the degree to which foreign knowledge is accessible and may vary due to the cultural and political relations between the domestic economy and the rest of the world.

The domestic planner's problem is allocating workers and R&D resources to maximize domestic welfare, taking the time path of import prices $\left\{p_t^f\right\}$ and foreign knowledge $\left\{\boldsymbol{q}_t^f\right\}$ as given:

$$V^*\left(\boldsymbol{q}_0, \left\{\boldsymbol{q}_t^f, p_t^f\right\}_{t=0}^\infty\right) \equiv \max_{\{s_{it}, \ell_{it}\}} \int_0^\infty e^{-\rho t} \ln \mathcal{C}\left(c_t^d, c_t^f\right) \mathrm{d}t, \tag{20}$$

subject to the open economy innovation production function (19); trade balance (18); goods production functions (2) and (3); the law of motion for domestic knowledge (5); and resource constraints (6). To ensure the planning problem (20) is well-defined, we assume $\left|\frac{\mathrm{d}\ln q_{it}^f}{\mathrm{d}t}\right|$ is bounded, and $p_t^f > 0$ is bounded away from zero.

Remark. We make three remarks on the economic environment with foreign spillovers. First,

while we assume x_{ij} to be exogenous and enters the innovation production function (19) log-linearly, our analysis holds as a first-order approximation in a richer environment where x_{ij} depends endogenously on the relative levels of domestic and foreign knowledge stock. This result, derived in Online Appendix B.7, is analogous to our closed-economy analysis in Section 2.5.

Second, we model international trade through the export of domestic bundle in exchange for the foreign bundle. This formulation significantly simplifies the exposition by removing the planner's incentive to use R&D allocation to manipulate the terms-of-trade (i.e., the relative prices of domestic products) and market power vis-a-vis foreign consumers.

Third, it is worth emphasizing that we do not analyze the optimal R&D allocation from the perspective of maximizing "global welfare". Doing so requires setting up a multi-country environment, which is certainly interesting and creates the opportunity to analyze a new set of other important issues—such as understanding cross-country R&D comparative advantages and strategic R&D policy—but it is beyond the scope of this paper.

2.6.2 Optimal R&D Allocation and Welfare with Knowledge Spillovers from Abroad

We now derive, in the setting with foreign spillovers, the optimal R&D allocation and the welfare impact of reallocating R&D resources.

Proposition 6. *Given paths of foreign knowledge and relative import prices* $\left\{\boldsymbol{q}_t^f, p_t^f\right\}_{t=0}^{\infty}$*, the solution to the open economy planner's problem (20) is time invariant and follows, along the entire time path,* $\ell_i/\bar{\ell} = \beta_i$ *and* $s_i/\bar{s} = \gamma_i$*, where*

$$\boldsymbol{\gamma}' = \xi^{-1}\frac{\rho}{\rho+\lambda}\boldsymbol{\beta}'\left(\boldsymbol{I} - \frac{\boldsymbol{\Omega}\circ\boldsymbol{X}}{1+\rho/\lambda}\right)^{-1}. \tag{21}$$

$\xi \equiv \frac{\rho}{\rho+\lambda}\boldsymbol{\beta}'\left(\boldsymbol{I} - \frac{\boldsymbol{\Omega}\circ\boldsymbol{X}}{1+\rho/\lambda}\right)^{-1}\mathbf{1}$ *is a constant that ensures* $\sum_i \gamma_i = 1$*;* $\boldsymbol{X} \equiv [x_{ij}]$ *is the matrix encoding the share of domestic contribution to cross-sector knowledge spillovers;* $\circ$ *is the Hadamard product.*

Proposition 6 generalizes Proposition 1 in Section 2.2 to an open economy. The ij-th entry of the Leontief inverse $\left(\boldsymbol{I} - \frac{\boldsymbol{\Omega}\circ\boldsymbol{X}}{1+\rho/\lambda}\right)^{-1} \equiv \sum_{m=0}^{\infty}\left(\frac{\boldsymbol{\Omega}\circ\boldsymbol{X}}{1+\rho/\lambda}\right)^{m}$ summarizes the network spillover effects from additional *domestic* knowledge in sector j on subsequent *domestic* innovation in sector i. Unlike in the closed economy, each round of network effect is no longer captured by the innovation network $\boldsymbol{\Omega}$ but is instead captured by $\boldsymbol{\Omega}\circ\boldsymbol{X}$: in the presence of knowledge spillovers from abroad, domestic R&D only contributes partially to the total knowledge spillovers from sector j to sector i; the elasticity of innovation efficiency in sector i with respect to domestic knowledge in sector j is captured by the ij-th entry of $\boldsymbol{\Omega}\circ\boldsymbol{X}$ (i.e., $\frac{\partial \ln \chi_i}{\partial \ln q_j} = \omega_{ij}x_{ij}$).

Proposition 6 highlights that countries with more self-contained innovation networks—such the U.S. and Japan, as we show later, where R&D builds more heavily on domestic rather than

on foreign knowledge—should allocate more R&D to network-central sectors. Conversely, countries that benefit more from foreign spillovers should direct R&D toward sectors that account for greater domestic value-added. Using our intuition from the closed economy Proposition 1, it is as if economies with self-contained innovation networks have patient planners while economies reliant on foreign knowledge have impatient planners. To see this, consider an economy in which the domestic share of knowledge spillovers is constant across all sector-pairs, $x_{ij} = x$. The Leontief inverse in (21) simplifies to $\left(\boldsymbol{I} - x \cdot \frac{\boldsymbol{\Omega}}{1+\rho/\lambda}\right)^{-1}$. Greater reliance on foreign knowledge (lower x) is therefore isomorphic to a higher discount rate ρ in a closed-economy, as if the planner place less value on long-term innovation spillovers.

The proportionality constant ξ in equation (21) ensures γ_i sums to one. It is a measure of R&D self-sufficiency. $\xi \leq 1$ in open economies and is decreasing in foreign dependence; $\xi = 1$ only if the economy does not benefit from any foreign spillovers ($x_{ij} = 1$ for all i, j).

Our next result provides the consumption-equivalent welfare impact of adopting the optimal R&D allocation, extending our closed-economy result in Proposition 4.

Proposition 7. *Consider an open economy with R&D self-sufficiency measure ξ and given paths of foreign knowledge and relative import prices $\left\{\boldsymbol{q}_t^f, p_t^f\right\}_{t=0}^{\infty}$. For any path of worker $\{\boldsymbol{\ell}_t\}$ and time-invariant R&D allocation $\boldsymbol{b}$, the consumption-equivalent welfare impact of adopting the optimal R&D allocation is*

$$\mathcal{L}(\boldsymbol{b}; \xi) = \exp\left(\psi\xi\frac{\lambda}{\rho}\boldsymbol{\gamma}'\left(\ln\boldsymbol{\gamma} - \ln\boldsymbol{b}\right)\right). \tag{22}$$

Relative to the closed-economy counterpart (15), the open-economy formula (22) depends additionally on ξ, the R&D self-sufficiency measure defined in Proposition 6. This term formalizes the notion that, in economies where domestic R&D matters less for long-run spillovers (lower ξ), suboptimally allocating domestic R&D is also less consequential for welfare.

2.7 Additional Results and Theoretical Extensions

We now discuss several additional theoretical results and extensions. Section 2.7.1 incorporates a production network into the baseline model. Section 2.7.2 embeds an innovation network into a semi-endogenous growth setting. Section 2.7.3 discusses potential inefficiencies in a decentralized equilibrium. Section 2.7.4 briefly describes several other extensions in the Online Appendix.

2.7.1 Production Network

Our baseline model features a simple production structure where all goods are produced directly from labor. We now discuss how to tractably incorporate a canonical production network into our framework. More details of this extension are provided in Section B.2 of the Online Appendix.

Conceptually, for the optimal R&D allocation $\boldsymbol{\gamma}' \propto \boldsymbol{\beta}' \left(\boldsymbol{I} - \frac{\boldsymbol{\Omega}}{1+\rho/\lambda} \right)^{-1}$, the presence of a production network requires a different construction for the $\boldsymbol{\beta}$ vector, but the innovation network $\boldsymbol{\Omega}$ term is unaffected. Formally, the $\boldsymbol{\beta}$ vector should capture the elasticity of aggregate consumption with respect to the productivity in each sector; in the presence of a production network, it should reflect not only the consumer preferences but also the production network structure. With this adjustment, our main results continue to hold in this environment.

Specifically, suppose the production of good i requires other goods as intermediate inputs, as in the canonical production network model (Acemoglu et al., 2012):

$$y_{it} = \left(q_{it}^{\psi} \ell_{it} \right)^{\alpha_i} \prod_{j=1}^{K} m_{ijt}^{\sigma_{ij}}, \quad \alpha_i + \sum_{j=1}^{K} \sigma_{ij} = 1, \tag{23}$$

where m_{ijt} is the quantity of good j used for the production of good i, α_i is sector i's output elasticity to value-added, and σ_{ij} is sector i's output elasticity to input j. The baseline model is a special case with $\sigma_{ij} = 0$ for all i, j.

Unlike in our baseline model, the elasticity of the aggregate consumption with respect to sectoral productivity ($\frac{\partial \ln y_t}{\partial \ln q_{it}^{\psi}}$) is no longer the consumption elasticity β_i; instead, standard results from the production network literature (e.g., see Hulten, 1978) imply that $\frac{\partial \ln y_t}{\partial \ln q_{it}^{\psi}} = \hat{\beta}_i \equiv \alpha_i \delta_i$, where $\boldsymbol{\delta}' = \boldsymbol{\beta}' \left(\boldsymbol{I} - \boldsymbol{\Sigma} \right)^{-1}$ is the influence vector, and $\boldsymbol{\Sigma} \equiv [\sigma_{ij}]$ is the input-output elasticity matrix. If the marginal product of labor is equalized across sectors, then $\hat{\beta}_i$ is the share of labor allocated to sector i. Our main results from previous sections, including the optimal R&D allocation and the welfare sufficient statistic, continue to hold when we use the vector $\hat{\boldsymbol{\beta}}$ in place of $\boldsymbol{\beta}$. In our empirical exercises, we measure this vector using sectoral value-added as a share of GDP, thereby accounting for the production network.

2.7.2 Semi-Endogenous Growth

Our baseline model features endogenous growth: a positive growth rate of aggregate output along a balanced growth path in the absence of population growth. This is because the R&D technology features aggregate constant-returns-to-scale in sectoral knowledge stock. To embed our innovation network formulation into a semi-endogenous growth setting with a constant growth rate in the total measure of scientists $\bar{s}_t = \bar{s}_0 e^{\bar{g} t}$, consider replacing the knowledge stock evolution equation (5) with

$$\dot{q}_{it} / q_{it} = \lambda \ln \left(n_{it} / q_{it}^{1+\kappa} \right),$$

where $\kappa \geq 0$ captures the rate at which proportional improvements in knowledge are getting harder to find (Bloom et al. 2020, Jones 2022). The knowledge law of motion (9) becomes

$$\mathrm{d} \ln \boldsymbol{q}_t \big/ \mathrm{d}t = \lambda \cdot \left(\ln \boldsymbol{\eta} + \ln \bar{s}_0 + \bar{g} t + \ln \boldsymbol{\gamma}_t + \left(\boldsymbol{\Omega} - (1+\kappa) \boldsymbol{I} \right) \ln \boldsymbol{q}_t \right).$$

Our baseline model corresponds to the special case where $\kappa = \bar{g} = 0$. When $\kappa > 0$, the R&D technology features aggregate decreasing-returns-to-scale, and the long-run growth rate of the economy is fully determined by the growth rate $\bar{g}$ of scientists. In Online Appendix B.3, we show that the optimal R&D allocation follows $\boldsymbol{\gamma}' \propto \boldsymbol{\beta}' \left(\boldsymbol{I} - \frac{\boldsymbol{\Omega}}{1+\kappa+\rho/\lambda} \right)^{-1}$, and the consumption-equivalent welfare impact of adopting the optimal allocation is $\mathcal{L}(\boldsymbol{b}) = \exp\left(\frac{\psi\lambda}{\rho+\kappa\lambda} \boldsymbol{\gamma}' (\ln \boldsymbol{\gamma} - \ln \boldsymbol{b}) \right)$. Compared to the formulas under endogenous growth, the semi-endogenous-growth version replaces the discount rate ρ with $\rho + \kappa\lambda$. Intuitively, in the presence of aggregate decreasing returns, growth in knowledge stock raises the difficulty of future R&D; from a welfare perspective, it is *as if* future growth are discounted at a higher rate.

2.7.3 Potential Inefficiency in A Decentralized Market

Why may a decentralized market not allocate R&D resources efficiently? In an innovation network, knowledge is a public good, as knowledge creation benefits subsequent R&D in other sectors and all future periods. To the extent that innovators do not fully internalize such future benefits,[3] a decentralized market does not implement the efficient R&D allocation.

To demonstrate the potential inefficiency, Online Appendix B.4 constructs a decentralized equilibrium in which innovators conduct R&D only in pursuit of profits—each innovation is patented, thereby granting the innovator a temporary stream of profits until replaced by a future innovation—disregarding any beneficial spillovers their R&D activities may provide in the future. As we show, the decentralized allocation of R&D resources along a BGP follows the consumption elasticity $\boldsymbol{\beta}$, which coincides with the planner's solution only if the society is completely myopic ($\rho/\lambda \to 0$).

This illustrative decentralized equilibrium lacks many real-world features of the market for innovation (e.g., multi-sector firms, mergers and acquisitions, and patent licensing). Nevertheless, it is important to note that, by comparing the R&D allocation in the data to the first-best, our notion of allocative efficiency—measured by the consumption-equivalent welfare impact of reallocating R&D optimally—does not require that we take a stance on firms' equilibrium conduct; instead, it directly calculates the welfare impact of reallocating R&D based on the economic environment specified in Section 2.1.

2.7.4 Other Theoretical Extensions

The Online Appendix includes a number of additional extensions that generalize our main results to various economic environments.

[3] While a patent gives its owner the legal right to exclude others from making or using an invention, it does not by design preclude others from benefiting from knowledge spillovers.

Section B.5 shows that our results extend naturally to a setting where the planner is constrained and can reallocate resources across only a subset of sectors: the constrained-optimal resource allocation within the subset is proportional to the unconstrained optimal allocation γ.

The open economy analysis in Section 2.6 considers a domestic planner who takes the paths of foreign knowledge as given. Section B.6 considers a domestic planner who internalizes the impact of domestic allocations on foreign variables. Section B.7 concerns a setting where both the innovation network (Ω) and the elasticities of foreign spillovers ($\boldsymbol{X}$) are endogenous.

Section B.8 considers an environment with sector-specific λ_i, the sensitivity of knowledge growth to the flow of new innovation relative to the existing stock.

Section B.9 considers a setting where the row-sums of Ω may differ (and not equal to one).

Section B.10 allows for factor mobility between production and R&D.

3 Data

This section describes the data for our empirical analyses. We use patent citation data across sectors and countries to construct the global innovation network. We also use data on sectoral production, final use, and R&D. Here we briefly describe how we construct and harmonize these data. Section C of the Online Appendix provides more details.

3.1 Data on Patents

U.S. Patents U.S. patent data are from the United States Patent and Trademark Office (USPTO).[4] This database provides detailed patent-level records for nearly seven million patents granted by the USPTO since 1976. The data include, for each patent, the application and grant years, the technology classifications based on the International Patent Classification (IPC) system, and the geographic locations of the patent assignee and patent inventors (the former holds legal ownership rights to the patent while the latter may not). Central to our network analysis, we observe each patent's citations of prior patents as well as the citations it receives from subsequent patents up to 2020, the year we extracted the data. In our empirical analysis, we use patents filed before the end of 2014 to mitigate the right-truncation problem, since patents filed more recently may still be in the approval process.

Global Patents To capture global innovation, we use Google Patents' global patent data, which contain information on more than 36 million patents from over 40 major patent authorities around the world, including those in the U.S., the European Union, Japan, and China, among others,

[4]We accessed the patent data via the USPTO PatentsView platform at https://www.patentsview.org/download/.

starting from the 18th century (data prior to the 1970s have limited coverage). Google Patents global innovation data are constructed based on the raw data records at DOCDB (EPO worldwide bibliographic data), which are the same records underlying other global patent datasets such as the PATSTAT database. As a result, Google Patents's data coverage and variable quality are nearly identical to those of PATSTAT. We choose to use Google Patents data for our main analysis because they are free of charge for any researcher, and we compare Google Patents to PATSTAT in detail in the Online Appendix D.

For each patent, Google Patents supplies information similar to the USPTO data described above. We assign each patent to a geographical unit using the country of residence of its inventors(s), country of residence of the assignee(s), and country of the patent authority, in that order. When a patent is associated with inventors or assignees from multiple countries, we attribute the patent to these countries assuming fractional and equal weight per assignee/inventor.

A major challenge when working with international patent data is multi-filing: to protect intellectual properties, it is common practice for innovators to file the same invention with multiple patent authorities in different countries, forming what is called a "patent family." To avoid double counting these inventions, our analysis uses only the first patent filed in each family when counting new innovation, while attributing all citations made to a whole family to this first patent. To identify patent family, we use the patent family ID assigned by Google Patents, self-reported multi-filing status, and the unique identifier for patents filed under the Patent Cooperation Treaty, which is an international law treaty aimed at protecting innovations across countries.

We measure the number of patents produced in a country-sector-year, both the raw counts and with quality adjustments using the number of citations each patent received. To capture actual patent timing, we use the year a patent was filed rather than granted.

3.2 Data on Production, Final Use, and R&D Allocation

Production and Final Use In our cross-country analysis, for each country and sector, we use the World Input-Output Database (WIOD, Timmer et al., 2015) to extract sectoral information. The data cover the years 2000–2014 and 43 major economies, which altogether represent more than 85% of world GDP. WIOD's sectoral categorization follows the two-digit International Standard Classification (ISIC) revision 4, with a total of 56 sectors covering the entire production spectrum, including primary, manufacturing, and service sectors. We obtain data on value-added (gross value-added, "VA"), employment ("empe"), output (gross output, "GO"). We also obtain information on intermediate inputs, value used for consumption, imports, and exports. For the U.S., we also obtain more detailed sectoral production, consumption, and import-export data, comprising 181 sectors from 1990 to 2019, from the Bureau of Labor Statistics (BLS).

Sectoral R&D Allocation Our quantitative analysis uses data on R&D allocation across different technology classes in each country. There is no standard database to exhaustively measure such information. Our primary measure relies on aggregating firm-level R&D expenditures to the country-sector-year level, based on three widely used firm-level data sets: Compustat, Worldscope, and Datastream. Combined, these data cover more than 110,000 global firms located in 160 countries and account for over 95% of the world's total market capitalization. For multinationals, we first attribute the firm-level R&D expenditures to IPC-country level in proportion to each firm's shares of patents in each IPC-country, following Griffith, Harrison, and Van Reenen (2006), and then aggregate to IPC-country-year level.

Our primary measure of sectoral R&D is imperfect, as the firm-level data sets oversample large firms and have potentially different reporting standard across countries; we also miss R&D inputs from public sectors. Nevertheless, it is important to note that, as our theory concerns the cross-sector R&D allocation, what matters for our quantitative analysis later is the allocation shares of R&D resources across sectors in each country and not the aggregate R&D levels; any mismeasurement that affects all sectors proportionally should have no quantitative impacts. As robustness checks, we later show that our primary measure of R&D allocation shares correlates strongly with two independent sources of R&D data, thereby giving us confidence in using our measure for quantitative analysis. The first robustness check calculates cross-sector R&D allocation using the innovation output (which is better measured) rather than input: the number of patents produced in each country-IPC (or country-sector) divided by total number of patents produced in that specific country (correlation with our primary measure of sectoral R&D averages to 0.74 across countries; see Appendix Table A.1). The second robustness check utilizes the OECD Analytical Business Enterprise Research and Development (ANBERD) Database (Machin and Van Reenen, 1998), which has country-sector-level R&D information. Relative to our primary R&D measure, the ANBERD Database has more limited country-year coverage and relies more on imputations from firm-level surveys. Our primary R&D measure also allows us to explicitly and transparently attribute R&D of multi-sector or multinational enterprises to different sectors and countries. Nevertheless, for all the major economies in both data sets, R&D allocation from ANBERD is highly correlated with our primary measure (see Appendix Table A.1), and our quantitative results are consistent using both data.

3.3 Concordance

Patent data are classified according to the IPC system, which is distinct from the classifications in our sectoral data. We build concordance between these two data types to construct sectoral measures of innovation activities and reversely to project sectoral measures into technology class

levels. To project patents from IPC onto sectors, we leverage the sectoral classifications covered in the three firm-level data sets described above. For the U.S., we link the USPTO patent database to Compustat using the bridge file provided by Kogan et al. (2017) and Ma (2021). Firms are classified by the North American Industry Classification System (NAICS) codes, which are then mapped to BLS sectors using the crosswalk file provided by the BLS website.[5] For the global analysis, we follow the analogous procedure and match Google Patents with global firm data from the Worldscope and Datastream databases. This process provides each patenting firm's International Standard Industrial Classification (ISIC), which can then be accurately mapped to WIOD that is also organized using the ISIC system. To reverse-project country-sector-year level measures onto country-IPC-year, we use the sector-IPC mapping provided in Lybbert and Zolas (2014).

We provide details on these matching procedures and the representativeness of using innovation measures aggregated from firm-level data in Section C of our Online Appendix.

4 Innovation Network and Knowledge Spillovers

In this section, we build several key data elements that will be used in our main quantitative analysis in Section 5. We first construct the innovation network $\mathbf{\Omega}$ and discuss its empirical properties. We then empirically validate a key mechanism in our model, that knowledge spillovers occur through innovation networks both domestically in the U.S. and globally.

4.1 Innovation Network

Constructing the Innovation Network We construct the innovation network from patent citations. First, we build the sector-to-sector innovation network. Let $Cites_{ijt}$ denote the total number of times that patents in sector i cite patents in sector j, among all patents filed in year t. As a baseline construction, we follow (Acemoglu et al., 2016) and define ω_{ijt} as the share of total citations patents in sector i made to sector j in year t:

$$\omega_{ijt} \equiv \frac{Cites_{ijt}}{\sum_{k=1}^{K} Cites_{ikt}}. \tag{24}$$

The object ω_{ijt} measures the extent to which upstream sector j's prior knowledge benefits innovation in sector i. The matrix $\mathbf{\Omega}_t$, whose ij-th entry is ω_{ijt}, captures the knowledge flow network we refer to as the *innovation network*. In a global setting, all subscripts will include additional country dimensions indicating the countries of sectors i and j, and the innovation network will measure the extent to which upstream country's sector j's prior knowledge benefits innovation in focal country's sector i. We can construct a country-specific network using patents from each

[5]The crosswalk can be accessed at: https://www.bls.gov/emp/documentation/crosswalks.htm.

country; we can also include patents from a time window broader than one year, such as using all patents over five or ten years.

The sector-to-sector innovation networks appear to be stable over time and across countries, suggesting their ability to capture some deep relations between sectoral innovation and technology diffusion. Table A.5 of the Online Appendix shows the serial correlation of entries in $\mathbf{\Omega}_t$ is near perfect a decade apart and remains above 0.8 even three decades apart. Table A.6 of the Online Appendix demonstrates that the innovation network constructed by pooling patents from all countries almost perfectly correlates with the U.S.-specific network (correlation 0.97) and highly correlates (correlation ≈ 0.8) with country-specific innovation networks from Japan, China, Germany, Canada, the United Kingdom, and France. Such high correlations imply that decisions about country and time specificities of the innovation network do not affect our analysis in major ways.

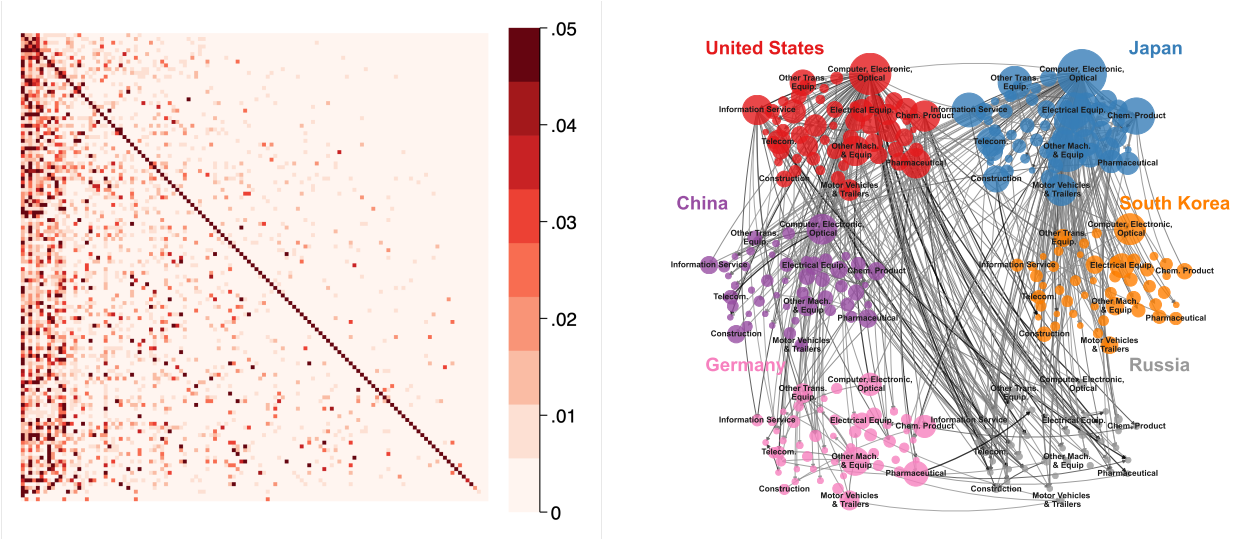
We should point out that the knowledge spillover network is inherently difficult to measure. Despite the evidence (Section 4.2 below in particular) for the usefulness of our citation-based construction, it is not perfect. In later analysis, we provide several alternative constructions of the network, such as weighting connections by the quality (total citations) of the cited or citing patents, to focus only on spillovers among major patents; we also create an innovation network $\mathbf{\Omega}$ with entries proportional to citations, without the normalization by each row's sum. Our quantitative findings are robust across all of these specifications.

Visualizing the Innovation Network Figure 1, Panel (a) visualizes the innovation network by plotting the matrix $\mathbf{\Omega}$ of 2010 as a heatmap. Each row and each column is a 3-digit IPC class, where the color in the i-th row and j-th column corresponds to ω_{ij} using the colormap listed to the right of the figure. Sectors are sorted by decreasing innovation centrality, the empirical properties of which we will formally discuss below. A key feature is that IPC classes follow a hierarchical structure: the innovation network is highly asymmetric, and there is a "pecking order" across sectors. Innovation-central sectors account for a disproportionate share of citations from all other sectors (columns are dense on the left but become progressively sparser to the right), yet these innovation-central sectors do not significantly cite noncentral sectors (rows are sparse on the top but become progressively denser toward the bottom).

Figure 1, Panel (b) visualizes the global innovation network by plotting each country-sector as a node, with size drawn in proportion to the total patent counts in our sample. An arrow from country m sector j to country n sector i indicates knowledge flow from mj to ni, with arrow width drawn in proportion to the share of ni's citations to mj. For visual clarity, only the largest countries and sectors are shown. Several patterns emerge from this figure. First, Japan and the U.S. produce the most patents in our sample. Second, the U.S. receives significantly more foreign

Figure 1. Visualizing the Innovation Network

(a) IPC-to-IPC (131×131) network $\boldsymbol{\Omega}$ **(b)** The global innovation network



Notes. The left panel visualizes the IPC-to-IPC (3-digit level) network $\boldsymbol{\Omega}$ as a heatmap, with darker colors representing larger matrix entries; sectors are ordered according to their innovation centrality. The right panel visualizes the global innovation network for six economies with the highest total patent output in our sample. Each node is a country-sector, with size drawn in proportion to patent output. Arrows represent knowledge flows, with width drawn in proportion to citation shares.

citations than any other economy in our sample; it is a major knowledge exporter and only a minor knowledge importer.

The Innovation Network Weakly Correlates with Input-Output Networks The innovation network $\boldsymbol{\Omega}$ encodes cross-sector linkages via knowledge spillovers. Another prominent type of cross-sector linkages occur through input-output relations, as sectors purchase intermediate inputs from one another during production. Table 1 shows that innovation and production networks are only weakly correlated. In other words, the two network relations capture different connections across sectors. Specifically, for each of the top ten countries ranked by total patent output, we compute the industry-by-industry input-output expenditure share matrix, which is a row stochastic matrix (as is $\boldsymbol{\Omega}$) commonly used to represent input-output relationships. Table 1 presents the correlation between entries in $\boldsymbol{\Omega}$ and those in the input-output matrix. The correlation is weak (<0.35) in all economies.

Innovation Centrality Across Sectors We provide some descriptive statistics of the innovation centrality $\boldsymbol{a}$, which is the dominant left eigenvector of the innovation network $\boldsymbol{\Omega}$. Recall that

Table 1. Correlations Between Country-Level Innovation Network and Production Network

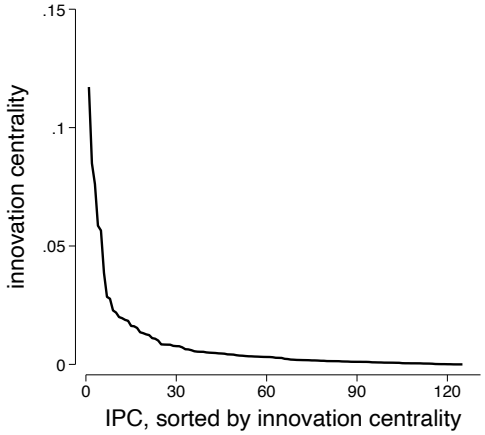
US	Japan	China	South Korea	Germany	Russia	France	UK	Canada	Netherlands
0.32	0.28	0.35	0.31	0.23	0.19	0.36	0.41	0.29	0.22

Notes. This table presents the correlations between the country-level innovation network matrix and the country-level input-output expenditure share matrix for the top 10 countries ranked by total patent counts during 2010–2014.

in our model, $\boldsymbol{a}$ is also the R&D allocation vector that maximizes the growth rate of a closed economy (Corollary 1) and is an important determinant of the optimal R&D allocation. The left panel of Figure 2 plots the innovation centrality a_i across 3-digit IPC sectors using the 2010 U.S. innovation network, where sectors are ordered along the x-axis in descending a_i. The figure shows innovation centrality is highly heterogeneous across sectors. To maximize economic growth, the most innovation-central sector should be allocated about twice as many R&D resources as the 5th sector ranked by a_i, about ten times as many as the 20th sector, and about 30 times as many as the 50th sector. The right panel of Figure 2 identifies the top 10 IPC classes; these include several technological classes related to medical science, computing, semiconductors, and electric communication technologies, among others.

Figure 2. Innovation Centrality and Key Sectors

(a) Innovation Centrality Across IPCs



(b) Top 10 IPCs By Innovation Centrality a_i

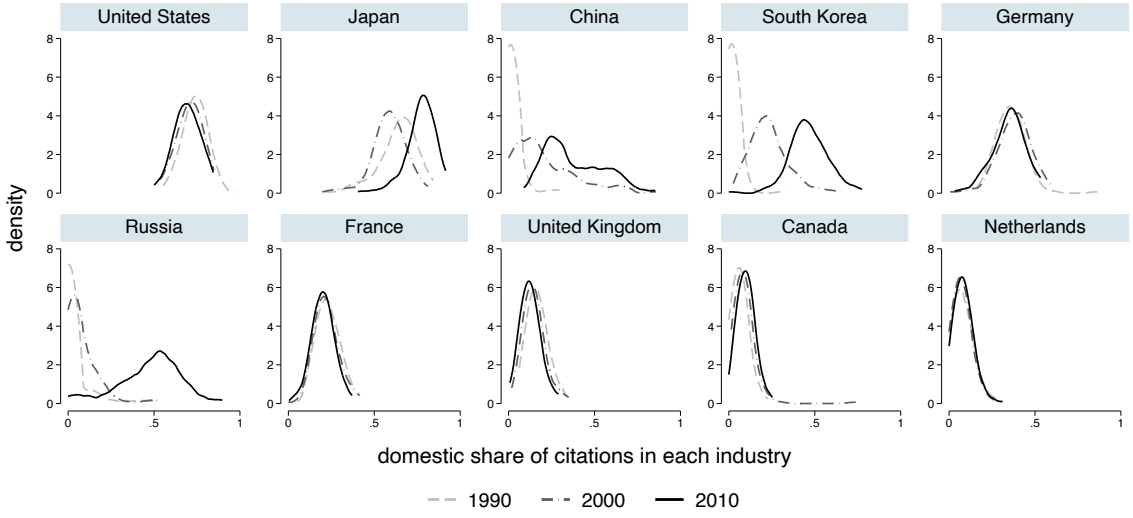
1	A61	medical or veterinary science; hygiene
2	G06	computing; calculating or counting
3	H01	basic electric elements
4	H04	electric communication technique
5	G01	measuring; testing
6	B60	vehicles in general
7	G02	optics
8	B01	physical or chemical processes or apparatus in general
9	C08	organic macromolecular compounds; their preparation or chemical working-up; compositions based thereon
10	F16	engineering elements or units; general measures for producing and maintaining effective functioning of machines or installations; thermal insulation in general

Notes. This figure presents the innovation centrality of different technology classes categorized using IPC. Panel (a) plots innovation centrality a_i across 3-digit IPC sectors ranked in descending order based on a_i. Panel (b) lists the top 10 IPC classes by their innovation centrality.

Cross-Country Linkages in the Innovation Network How much do countries benefit from foreign knowledge? To answer this, for each country m, sector i, and year t, we compute the domestic share of citations made by patents in country-sector mi. Figure 3 shows the distribu-

tion of domestic citation shares across all sectors for the ten economies with the highest patent counts in our sample, for the years 1990, 2000, and 2010. The U.S. relies relatively sparingly on foreign knowledge: consistently across these three decades, about 70% of citations by U.S. patents are made to other U.S. patents. In contrast, citations made to foreign patents account for the vast majority of citations by all other economies except Japan, suggesting these economies benefit significantly from foreign knowledge, most notably from the U.S. The Japanese self-citation shares increased over time on average, from 65% in 2000 to 77% in 2010. Declining foreign reliance over time is also observed for China and South Korea, although their levels of foreign reliance remain high.

Figure 3. Cross-Sector Distribution of Domestic Citation Shares by Country



Notes. This figure presents the distribution of each country's domestic citation shares across sectors, showing the distribution using 1990, 2000, and 2010 data. Sector definitions follow 3-digit IPC classes. Domestic citation share for each country-sector is defined as the number of citations made to domestic patents as a share of total citations made by new patents invented in that each country-sector.

4.2 Knowledge Spillovers Through the Innovation Network

We now provide evidence that knowledge spillovers occur through the innovation network, with the purpose of validating our model mechanism and the innovation network construction. We build on Acemoglu, Akcigit, and Kerr (2016)—which provide similar evidence for the U.S. domestic network—and extend using instrumental variables (IVs) and to the global setting. The IVs are constructed based on time-varying sectoral exposure to tax-induced user cost of R&D (Wilson, 2009 and Thomson, 2017); they help isolate movements in patent output driven by knowledge spillovers and not by common shocks to connected sectors (Manski, 1993 and Bloom, Schanker-

man, and Van Reenen, 2013). We find evidence for directional knowledge spillovers: each sector's innovation output responds only to past upstream innovations and does not respond to past innovation from downstream sectors even though they are also connected. We also show that relative to input-output linkages, the innovation network is a significantly stronger channel through which knowledge spillovers take place.

4.2.1 U.S. Evidence

We first test the mechanism in the U.S. using an empirical specification derived by our model, treating the U.S. as a closed economy. Specifically, integrating our law of motion (5) over time, we can express the knowledge stock as $\ln q_{jt} = \int_0^\infty e^{-\lambda s} \ln n_{j,t-s} \, \mathrm{d}s$. The innovation production function (4) further implies a log-linear relationship among sector i's new patents, sectoral R&D, and past patents from other upstream sectors:

$$\ln n_{it} = \ln \eta_i + \ln s_{it} + \lambda \underbrace{\sum_{j=1}^{K} \omega_{ij} \left(\int_0^\infty e^{-\lambda s} \ln n_{j,t-s} \, \mathrm{d}s \right)}_{\equiv \ln \chi_{it}, \text{ the aggregation of knowledge that benefits innovation in sector } i \text{ at time } t}. \tag{25}$$

Equation (25) implies that, after controlling for sectoral R&D expenditures, past patents $\ln n_{j,t-s}$ in sector j influence new patent output in sector i through the innovation network ω_{ij}. Equation (25) also shows that, importantly, knowledge spillovers' effect is directional: the knowledge flow from sector j to sector i operates through ω_{ij} and not ω_{ji}.

We test the discrete-time analogue of (25) by constructing the knowledge aggregator χ_{it} from past patents. As a baseline measure, for each sector i, we enumerate over all sectors j from which knowledge flows to i, aggregating j's log patent counts $\ln n_{j,t-\tau}$ in the past 10 years ($1 \leq \tau \leq 10$), weighted by $\omega_{ij,t-\tau}$, the share of citation made from i to j in the corresponding year:

$$\text{Knowledge}_{it}^{Up} \equiv \textstyle\sum_{j \neq i} \sum_{\tau=1}^{10} \omega_{ij,t-\tau} \ln n_{jt-\tau}. \tag{26}$$

$\text{Knowledge}_{it}^{Up}$ captures the stock of past knowledge "upstream" of sector i and is the empirical counterpart to $\ln \chi_{it}$. We then perform the following regression, with sector and year fixed effects:

$$\ln n_{it} = \beta_1 \times \text{Knowledge}_{it}^{Up} + \beta_2 \times \ln R\&D_{i,t-1} + \xi_i + \xi_t + \epsilon_{it}, \tag{27}$$

where n_{it} is the number of patents filed in sector i year t that are eventually granted. $R\&D_{i,t-1}$ is the R&D expenditure, with a one-year lag to reflect the delayed nature of patent filing (our results are robust to controlling for concurrent or additional lags of R&D expenditures).

Note that, when constructing the upstream knowledge aggregator (26), we exclude the lagged patent output from each sector itself. Doing so ensures the coefficient β_1 in regression (27) is not driven by serially correlated shocks to sectoral patent output, but our results are robust to includ-

ing lagged patent output from own sector as an additional regressor (Appendix Table A.8). Also note that, theoretically, the knowledge aggregator in (25) discounts past patents exponentially. Our empirical construction (26) features a discrete cutoff window ($\tau \leq 10$ years) to be agnostic about the discount factor λ, but our results are robust to alternative values of τ (see Appendix Table A.9 for $\tau = 5$ and 20), exponential discounting (Appendix Table A.10), or estimating the impact of past upstream patents nonparametrically at different time lags (Appendix Figure A.9).

Table 2. U.S. Evidence of Innovation Spillovers Through the Innovation Network

Y=	ln(Patents)				ln(Cites)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$Knowledge_{it}^{Up}$	0.586***	0.600***	0.508***	0.679**	0.802***	0.830***	0.743***	0.974***
	(0.180)	(0.205)	(0.174)	(0.266)	(0.202)	(0.218)	(0.196)	(0.279)
$ln(R\&D)_{i,t-1}$	0.275***	0.274***	0.279***	0.269***	0.258***	0.256***	0.261***	0.174**
	(0.063)	(0.062)	(0.060)	(0.070)	(0.086)	(0.086)	(0.086)	(0.082)
$Knowledge_{it}^{Down}$		-0.029				-0.058		
		(0.157)				(0.098)		
$Knowledge_{it}^{Up,IO}$			0.363**				0.268	
			(0.173)				(0.205)	
Specification	OLS	OLS	OLS	IV 2nd Stage	OLS	OLS	OLS	IV 2nd Stage
IV 1st Stage F-statistics				465.9				465.9
R^2	0.915	0.915	0.917	0.152	0.901	0.901	0.902	0.099
No. of Sectors	94	94	94	94	94	94	94	94
No. of Obs	1847	1847	1847	1113	1847	1847	1847	1113
Fixed Effects	Sector, Year				Sector, Year			

Notes. This table tests the relation between innovation in a focal sector and past innovation in sectors connected through the innovation network, using the U.S. data from BLS sectors. We restrict the sample to sectors that have at least 100 patents over the full sample period. Standard errors in parentheses are clustered at the sector level. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 2, column (1) presents the results of regression (27). Sectoral R&D expenditure significantly predicts the number of new patents filed in a given year, with an elasticity of 0.275. The knowledge stock upstream of each sector ($\text{Knowledge}_{it}^{Up}$) also significantly predicts patent output, with an elasticity of 0.586. Column (5) shows that both variables also predict patent quality: sectors with greater R&D and greater upstream knowledge stock tend to produce patents with more future citation counts. In the Online Appendix Table A.11, we also demonstrate these variables predict the commercial value of innovation measured using the stock market reaction upon patent approval (Kogan et al., 2017).

These regressions provide supportive evidence that past knowledge in sectors upstream of i benefits subsequent patent production in the focal sector i. An alternative story relates to common shocks: a group of sectors connected to each other via citation linkages may face similar demand, supply and investment opportunities, leading to co-movements of innovation activities. Serial correlations in such common shocks would lead to a positive coefficient β_1 in regression (27) even without cross-sector knowledge spillovers. This is a version of the "reflection problem"

à la Manski (1993) and Bloom et al. (2013).

We implement three additional analyses to address the "common shock" concern. First, we exploit the directional nature of knowledge spillovers. We construct the knowledge stock aggregator for sectors *downstream* of i in the innovation network:

$$\text{Knowledge}_{it}^{Down} \equiv \sum_{k \neq i} \sum_{\tau=1}^{10} \omega_{ki,t-\tau} \ln n_{kt-\tau}.$$

$\text{Knowledge}_{it}^{Down}$ aggregates the patent output in all sectors $k \neq i$, weighted by the extent to which patents in sector k cite those in sector i, and is therefore a measure of the knowledge stock *downstream* of sector i. Because knowledge flow is directional, there should be an asymmetry: while $\text{Knowledge}_{it}^{Up}$ should positively predict subsequent patent output in sector i, $\text{Knowledge}_{it}^{Down}$ should not. Yet any common shocks hitting this network should generate symmetric correlations in innovation output for focal sector i and both its upstream and downstream sectors.

Columns (2) and (6) of Table 2 add $\text{Knowledge}_{it}^{Down}$ to our baseline regressions. We make two observations. First, adding $\text{Knowledge}_{it}^{Down}$ as a control does not meaningfully affect the economic or statistical significance of our two baseline variables. This suggests our baseline regressions are not simply picking up correlated shocks to local technology clusters. Second, the coefficient on $\text{Knowledge}_{it}^{Down}$ is precisely zero, confirming our key model mechanism and that knowledge flow along the innovation network is directional—it goes only from upstream to downstream, and not the other way around.

Another related concern is that common shocks operate not through technological linkages but through input-output (IO) linkages. To address this, we construct the aggregator $\text{Knowledge}_{it}^{Up,IO}$ similarly to $\text{Knowledge}_{it}^{Up}$, but using patents from other sectors weighted not by the innovation network, as in (26), but instead by sector i's cost share on inputs from sector j. Columns (3) and (7) of Table 2 show the regression results when including $\text{Knowledge}_{it}^{Up,IO}$. Knowledge from innovation-upstream sectors remains an economically and statistically significant predictor of subsequent innovation in the focal sector. By contrast, $\text{Knowledge}_{it}^{Up,IO}$ has a smaller impact on sector i's patent quantity, and insignificant effect on innovation quality in these specifications. These results, along with the fact that the innovation network only weakly correlates with the IO network (see Table 1), imply that the innovation network provides valuable incremental information that is particularly powerful for understanding knowledge spillovers across sectors.

Next, we adopt another approach to address the "common shock" concern using tax-induced changes to the effective cost of R&D to create exogenous variations in innovation activities, following Bloom, Schankerman, and Van Reenen (2013). We briefly describe the approach here and provide more details in the Online Appendix E.3. The approach leverages the fact that the user-cost of R&D capital (i.e., the cost of conducting R&D) varies with state-level R&D tax credit, depreciation allowance, and corporate tax rate. Cross-sector heterogeneity in the geographic dis-

tribution of R&D activities in turn translates into R&D cost differences across sectors and over time. For the U.S., we use Wilson (2009)'s estimates of state-specific R&D cost shifters, combined with our estimates of the cross-state distribution of each sector's R&D, to calculate a sector's R&D costs. For the global setting below, we follow Thomson (2017) to calculate the user cost of R&D capital at the country-sector-year level.

We first create fitted values of sectoral innovation output ($\ln n_{it}$) using R&D cost shifters; we then use these fitted values in equation (26) to construct a predicted value of $\widehat{\text{Knowledge}}_{it}^{Up}$, which is in turn used as an instrumental variable (IV) for $\text{Knowledge}_{it}^{Up}$ for a two-stage least-squares (2SLS) analysis. Columns (4) and (8) in Table 2 remain qualitatively and quantitatively robust to using this IV strategy. Details of this analysis are provided in Online Appendix E.3.

4.2.2 Global Evidence

We now test international knowledge spillovers in our global sample. First, we construct an analogous measure of upstream knowledge stock: for each focal country m, sector i in year t, we enumerate over all countries c and sectors j in our sample, aggregating the (log-)patent output in cj over the past 10 years, weighted by the share of mi's citations that are to cj in the corresponding year:

$$\text{Knowledge}_{mit}^{Up} \equiv \sum_{cj \neq mi} \sum_{\tau=1}^{10} \frac{Cites_{mi \to cj, t-\tau}}{\sum_{c'=1}^{N} \sum_{k=1}^{K} Cites_{mi \to c'k, , t-\tau}} \ln n_{cj,t-\tau}. \tag{28}$$

Next, we adapt our closed economy test of knowledge spillovers to perform on the global innovation network. In this case, the unit of observation is at the country-industry-year level, and we include a saturated set (country-industry, country-year, industry-year) of fixed effects:

$$\ln n_{mit} = \beta_1 \times \text{Knowledge}_{mit}^{Up} + \beta_2 \times \ln R\&D_{mi,t-1} + \xi_{mi} + \xi_{mt} + \xi_{it} + \varepsilon_{ict}. \tag{29}$$

Table 3 shows the results: knowledge stock upstream of each country-industry significantly predicts subsequent patent counts (column 1) and citation-adjusted patent counts (column 5) even in the global setting. The coefficients are lower than estimates based only on the U.S., suggesting knowledge spillovers are stronger across sectors within the U.S. than they are across countries.

To rule out common shocks to technological and input-output clusters, we again make use of knowledge aggregated from downstream, from the input-output network, and the tax-induced IV, and find evidence in support of knowledge spillovers through the innovation network.

5 Application: R&D Resource Allocation in the Data

This section hosts our main empirical analysis, which uses our model to study the allocation of R&D resources in the data. We present the computed unilaterally optimal allocation of R&D

Table 3. Global Evidence of Knowledge Spillovers Through the Innovation Network

$Y=$	ln(Patents)				ln(Cites)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$Knowledge^{Up}_{mit}$	0.157***	0.199***	0.154***	0.202*	0.348***	0.424***	0.345***	0.405***
	(0.057)	(0.061)	(0.057)	(0.113)	(0.083)	(0.089)	(0.084)	(0.147)
$ln(R\&D)_{mi,t-1}$	0.033***	0.034***	0.032***	0.066***	0.046***	0.046***	0.046***	0.072***
	(0.010)	(0.010)	(0.010)	(0.014)	(0.014)	(0.014)	(0.014)	(0.022)
$Knowledge^{Down}_{mit}$		-0.091**				-0.167***		
		(0.043)				(0.059)		
$Knowledge^{Up,IO}_{mit}$			0.079				-0.038	
			(0.064)				(0.070)	
Specification	OLS	OLS	OLS	IV 2nd Stage	OLS	OLS	OLS	IV 2nd Stage
IV 1st Stage F-statistics				146.2				146.2
R^2	0.969	0.969	0.969	0.040	0.944	0.944	0.945	0.031
No. of Country x Sectors	564	564	550	280	564	564	550	280
No. of Obs	10,552	10,552	10,318	4,467	10,552	10,552	10,318	44,67
Fixed Effects	Country x Sector, Country x Year, Sector x Year				Country x Sector, Country x Year, Sector x Year			

Notes. This table tests the relation between innovation in a focal sector and past innovation in connected sectors through the innovation network, in an international setting. We restrict the sample to country-sectors with at least 10 patents over the full sample period. Standard errors in parentheses are clustered at the country-sector level. *, **, and *** indicate significance at the 10%, 5%, and 1% levels respectively.

resources across sectors for each country in our sample. We show that on average, sectors that should have more R&D resources do receive more resources, especially for the five economies with the most patents during our sample period. Nevertheless, the residual misalignment between the optimal and actual allocations remains large and is highly heterogeneous across countries. We compute the welfare gains from adopting the optimal R&D allocation for each country.

5.1 Optimal R&D Allocation

For each country and year, we calculate the unilaterally optimal cross-sector allocation of R&D resources $\boldsymbol{\gamma}$, using Proposition 6:

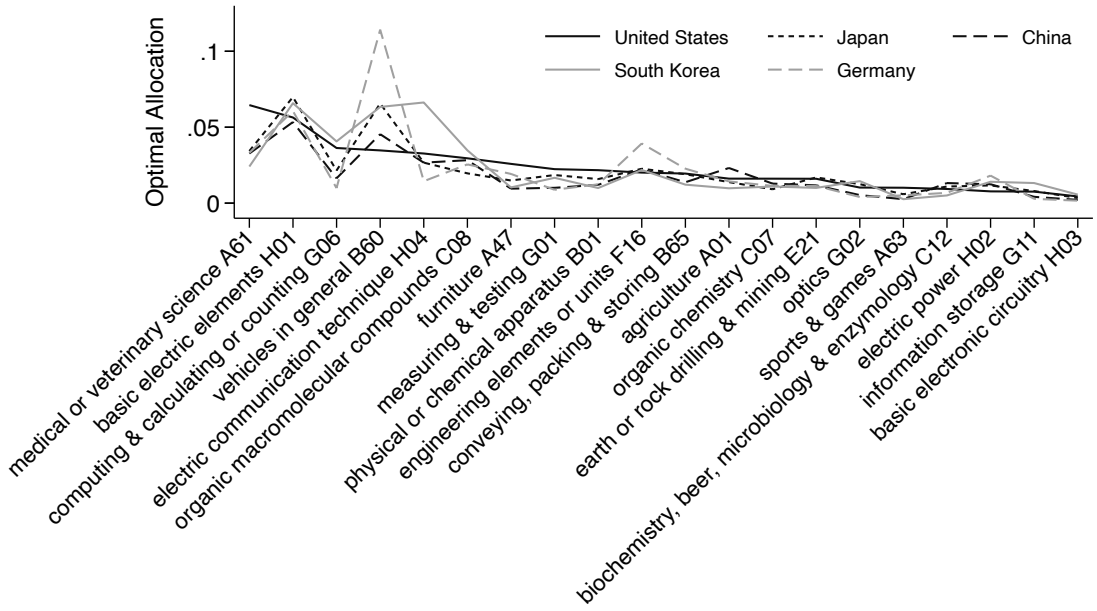
$$\boldsymbol{\gamma}' = \xi^{-1} \frac{\rho}{\rho + \lambda} \boldsymbol{\beta}' \left(\boldsymbol{I} - \frac{\boldsymbol{\Omega} \circ \boldsymbol{X}}{1 + \rho/\lambda} \right)^{-1}, \tag{30}$$

where the proportionality constant ξ ensures that elements in the optimal allocation vector $\boldsymbol{\gamma}$ sum to one and is a measure of R&D self-sufficiency. We measure $\boldsymbol{\beta}$ using each country's sectoral value-added relative to GDP in that year (thereby accounting for input-output linkages; see Section 2.7.1). Recall that $\boldsymbol{X} \equiv [x_{ij}]$ is the matrix encoding the share of domestic contribution to cross-sector knowledge spillovers; we measure x_{ij} as the share of citations from i to sector j that are toward domestic patents in j. As Figure 3 shows, entries of $\boldsymbol{X}$ average to above 70% across sectors for the U.S. but are significantly lower for all other countries except Japan in recent years.

To implement the formula in (30), we need to specify the discount rate ρ and the elasticity λ of knowledge growth with respect to the new innovation flows. As a baseline, we set discount

rate $\rho = 5\%$. Because knowledge stock does not have a natural scale, we rely on the empirical pattern of knowledge spillover dynamics to calibrate λ. Specifically, equation (25) implies that the spillover effect of past upstream knowledge on current innovation decays over time at rate λ. We estimate the spillover effects of past upstream patents nonparametrically at different time lags (Appendix Figure A.9). We find that the knowledge spillovers have a half-life of about four years, corresponding to $\lambda = 0.17$, which we adopt as the baseline calibration. Qualitatively, our analysis is not sensitive to the value of ρ/λ: as we show in the Online Appendix Table A.21, the optimal allocation $\boldsymbol{\gamma}$ is highly correlated across specifications with alternative values of ρ/λ.[6] The Online Appendix also reports additional results and sensitivity checks such as using data from other sample periods, using alternative specifications for $\boldsymbol{\Omega}$, and calibrate sector-specific λ_i. We describe these and other extensions in Section 5.4.

Figure 4. Optimal R&D Allocations in Different Countries



Notes. This figure shows the optimal R&D allocation across 20 3-digit IPC classes with the most patents for the five economies that produced the most patents in 2010–2014. Optimal allocations are calculated using our baseline calibration $\rho = 5\%$,$\lambda = 0.17$. Sectors are sorted by the optimal allocation for the U.S.,γ_{US}.

Figure 4 plots the optimal R&D allocation $\boldsymbol{\gamma}$ for the five economies that produced the most patents in 2010–2014. For visual clarity, we only show the top 20 3-digit IPC classes ranked by

[6]One reason why our conclusions are robust to alternative values of ρ/λ is that, as discussed previously, an increase in ρ/λ has the same implication for the optimal R&D allocation as an increase in a country's reliance on foreign knowledge. The substantial cross-country variation in foreign reliance (see Figure 3) dwarfs reasonable variation in the calibration of ρ/λ. Hence, the qualitative cross-country differences in unilaterally optimal R&D allocations are not sensitive to our calibration of ρ/λ.

total patent counts; these 20 classes account for 75% of all patents. The level of optimal R&D resources is shown on the y-axis, and the x-axis represents IPCs (sorted by γ_{US}).

For the U.S. (solid black line), the optimal allocation favors sectors with the highest innovation centrality, as listed in Figure 2, such as medical science (A61), basic electric elements (H01, e.g., semiconductors), and computing devices (G06). The top 10 IPCs (out of 131) should receive about a third of total U.S. R&D resources. The correlation between optimal U.S. R&D allocation and the innovation centrality $\boldsymbol{a}$ is 0.75. The correlation is high because the U.S. has a self-contained innovation network with relatively few citations toward foreign patents; hence, its planner should internalize more knowledge spillovers. The correlation is not perfect since the planner also considers IPC's importance for domestic production, encoded in the value-added share vector $\boldsymbol{\beta}$, which raises the optimal allocation of high-β sectors such as vehicles (B60).

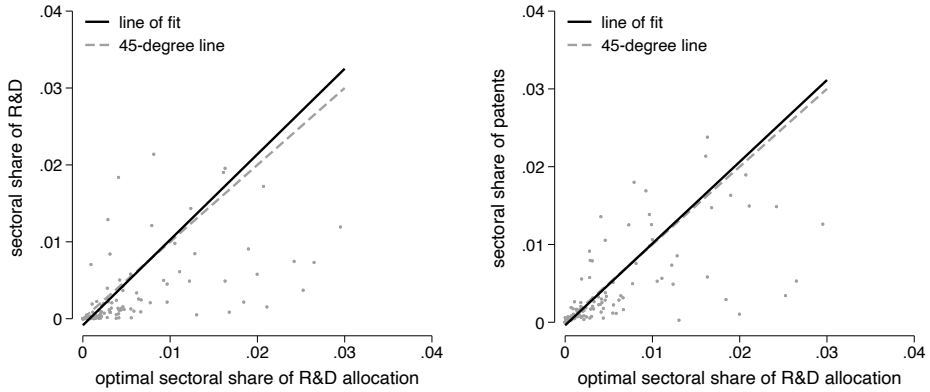
Figure 4 also reveals cross-country variations in the unilaterally optimal R&D allocations. Relative to the U.S., Germany and Japan should allocate more toward vehicles (B60); South Korea should allocate more toward electric communication technique (H04); all four non-U.S. economies should allocate less toward medical science (A61).

5.2 Innovation Allocation in the Data

We first present our model's ability to fit R&D resource allocation in the U.S. The left panel of Figure 5 shows the scatter plot of sectoral R&D expenditure (as a share of total R&D) against the optimal R&D expenditure share γ_{US} for the sample period 2010–2014. The linear fit (solid line) is close to the 45-degree line (dashed) with a slope of 1.11 (t-statistic 7.64), indicating that on average, sectors that should optimally receive more R&D resources do indeed receive more R&D resources. In the right panel of Figure 5, we change the y-axis to sectoral patent output as a share of total patent output; again, sectoral patent output aligns very well with γ_{US}, with a slope of 1.05 (t-statistic 8.20). To be clear, the strong alignment between real-world and optimal R&D allocations does not imply the U.S. allocates R&D optimally: there is substantial residual variation in R&D allocations as they disperse around the 45-degree line. The vertical distance between each observation and the 45-degree line measures the amount of R&D resources that need to be reassigned to achieve the optimal allocation. We quantitatively assess the welfare gains from adopting the optimal R&D allocation in Section 5.3 below.[7]

[7] A potential concern is that Figure 5 picks up a mechanical relationship: it may be that sectors with more resources produce more patents and citations, thereby appearing to be more central in the innovation network $\boldsymbol{\Omega}$—in other words, allocated resources reversely affect centrality of the innovation network. To argue against this possibility, we reproduce our empirical exercises using the innovation network constructed using citations from Japanese patents to Japanese patents. Because Japan's innovation network is self-contained and has few citations toward foreign patents (Figure 3), the network is by construction independent of U.S. R&D. All of our findings continue to hold, suggesting that innovation centrality $\boldsymbol{a}$—which correlates strongly with the U.S. optimal R&D allocation $\boldsymbol{\gamma}^{US}$—indeed picks up sectoral importance in the innovation network rather than representing historical R&D ex-

Figure 5. U.S. Actual R&D Allocation vs. Optimal Allocation$\boldsymbol{\gamma}_{US}$



Notes. This figure shows scatter plots of real-world sectoral R&D expenditure shares (left panel) and patent output shares (right panel) against optimal R&D allocation shares, γ_{US}, for the U.S. in 2010-2014. The solid line is the linear fit; the dashed line is the 45-degree line.

There is substantial cross-country heterogeneity in R&D resource allocations. Figure 6 shows scatter plots of sectoral R&D expenditure shares against the unilaterally optimal R&D allocations for ten countries that filed the most patents in 2010–2014. Sectoral R&D expenditure correlates strongly with the optimal R&D allocations for the five countries shown in the top row (U.S., Japan, China, South Korea, and Germany), and the relationship is weaker for the five economies at the bottom (Russia, France, U.K., Canada, and Netherlands). As we have noted, a line-of-fit with a slope of 1 does not imply resources are allocated optimally; nevertheless, Figure 5 suggests that on average, more resources need to be reallocated to achieve optimality in the five economies in the bottom row. These results are robust to using patent output share and R&D shares reported in OECD ANBERD database as measures of cross-sector R&D resource allocation (Figure A.12 and A.14 in the Online Appendix), suggesting that our results are not driven by the coverage and quality of our R&D expenditure variable. In Figure A.11 in the Online Appendix, we plot the alignment between real and optimal allocation in the years 2000 and 2005 and find very similar patterns.

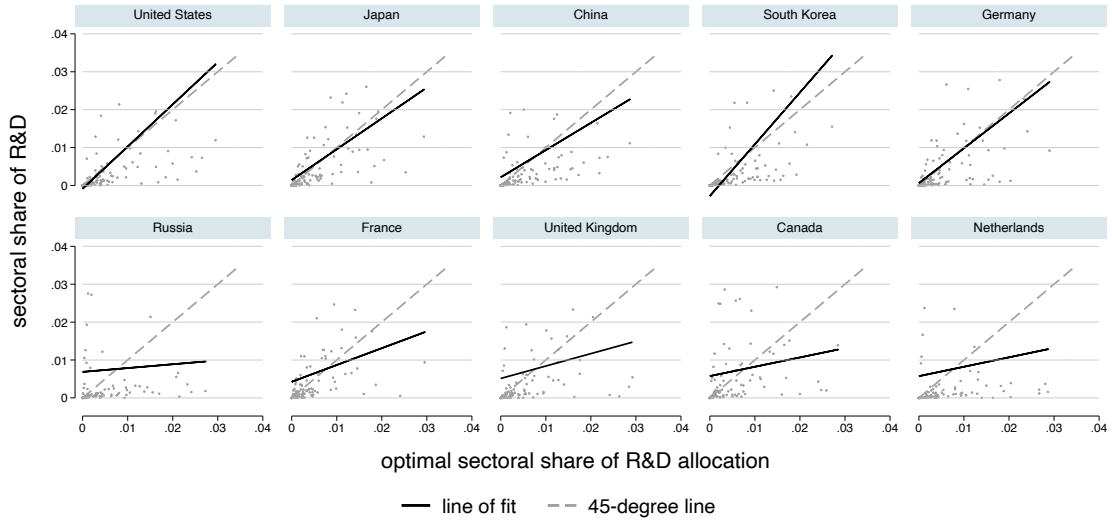
5.3 Welfare Gains from Improving R&D Allocation

We now quantify the potential welfare gains from improving R&D allocation, using the welfare formula in Proposition 7:

$$\ln \mathcal{L}\left(\boldsymbol{b}, \xi\right) = \psi \frac{\lambda}{\rho} \underbrace{\xi}_{\text{self sufficiency}} \times \underbrace{\boldsymbol{\gamma}'\left(\ln \boldsymbol{\gamma} - \ln \boldsymbol{b}\right)}_{\text{R\&D misallocation}}, \tag{31}$$

penditures.

Figure 6. Actual R&D Allocation vs. Optimal Allocation Across Countries



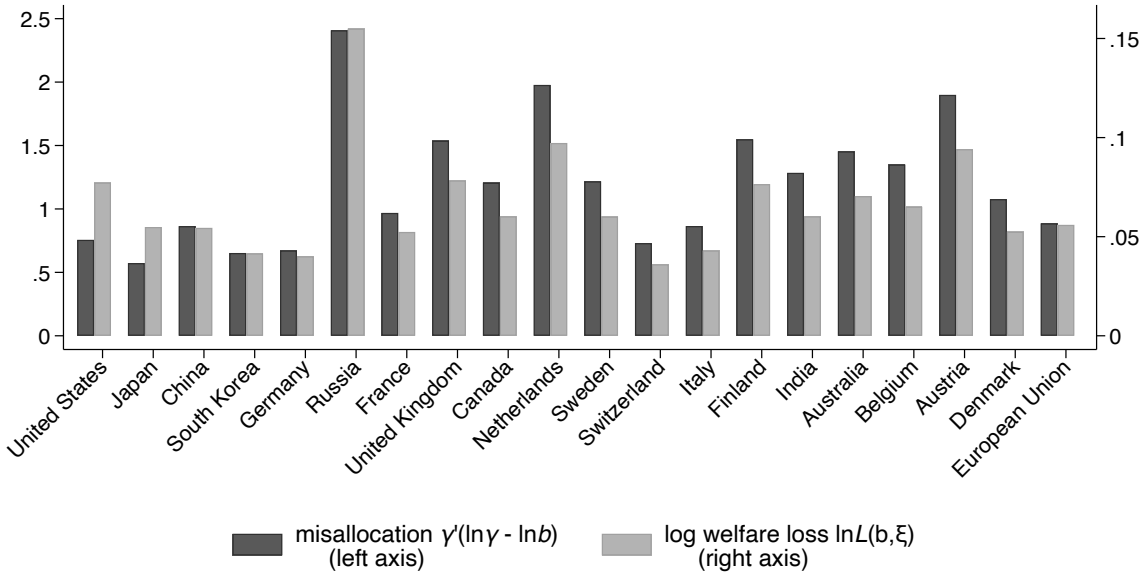
Notes. This figure shows scatter plots of sectoral patent output (as a fraction of total patent output in each country) against the optimal sectoral share of R&D allocation for the top ten innovative countries (in terms of patent output) using data from 2010–2014. The solid line is the linear fit; the dashed line is the 45-degree line. For visual clarity, outliers—sectors that account for >5% of national patent output—are not shown in the scatter plots, but all sectors are used when constructing the linear fit.

where $\ln \mathcal{L}(\boldsymbol{b}, \xi)$ is the consumption-equivalent welfare gains in logs, $\boldsymbol{b}$ is the empirical R&D allocation vector, $\boldsymbol{\gamma}$ is the optimal R&D allocation vector, and $\xi \equiv \frac{\rho}{\rho+\lambda}\boldsymbol{\beta}'\left(\boldsymbol{I} - \frac{\boldsymbol{\Omega} \circ \boldsymbol{X}}{1+\rho/\lambda}\right)^{-1}\mathbf{1}$ is the scalar measure of R&D self-sufficiency, which increases in domestic citation shares x_{ij}. ψ captures the elasticity of productivity to knowledge stock and thus proportionally controls the welfare impact of R&D allocation. It is an additional parameter to be calibrated. We specify $\psi = 0.06$, which implies a semi-elasticity $(\mathrm{d}g^y/\mathrm{d}\ln\bar{s} = \psi\lambda =)$ 0.01 of BGP consumption growth rate to the total stock of R&D resource, consistent with standard calibrations in the growth literature (Akcigit and Kerr, 2018, Akcigit et al., 2021).

Equation (31) implies a natural decomposition when comparing across countries the potential welfare gains from reallocating R&D: misallocation, measured as relative entropy from the actual allocation $\boldsymbol{b}$ to the optimal $\boldsymbol{\gamma}$ (higher $\boldsymbol{\gamma}'(\ln\boldsymbol{\gamma} - \ln\boldsymbol{b})$ means less efficient) and self-sufficiency (ξ). Holding an economy's R&D self-sufficiency constant, worse misallocation implies larger potential gains; and, as an economy benefits more from foreign knowledge spillovers (lower ξ), domestic R&D misallocation becomes less consequential for consumer welfare.

Figure 7 shows the degree of misallocation and the size of the potential welfare gains from R&D reallocation for the 19 economies that filed the most patents during 2010–2014. Because of the EU's high degree of economic integration, we also aggregate R&D from all EU countries and calculate the allocative efficiency of the EU as a single, integrated economy, which is listed

Figure 7. R&D Allocative Efficiency and Potential Welfare Gains Across Countries



Notes. This table shows the level of R&D misallocation (dark bars) and the potential welfare gains (light grey bars) from adoption optimal R&D allocation across 19 innovative countries with the highest patent outputs in our sample, and the integrated economy of the European Union, using 2010-2014 data. The calculation focuses on R&D in top 50 IPC classes by total patents.

as the 20th economy in Figure 7. The dark bars represent the misallocation term (left Y-axis). Among high-income countries, Japan has the most efficient R&D allocation, followed by other top-patenting countries like the U.S., France, and South Korea. Among high-income economies in Europe, Germany also has a higher allocative efficiency compared to its peers. When evaluated as a single integrated economy, the EU's allocative efficiency is comparable to Switzerland's and is marginally less than Germany's.

The light grey bars in Figure 7 represent the potential welfare gains of reallocating R&D optimally (right Y-axis). By our welfare accounting formula (31), the welfare gain is proportional to the misallocation term times the R&D self-sufficiency measure ξ. For the two economies with self-contained innovation networks, namely the U.S. and Japan, ξ is closer to one, so the overall welfare gains (grey bars) are closer in magnitude to the misallocation terms (dark bars). By contrast, the welfare gains are comparatively lower than the corresponding misallocation terms in all other economies, as their domestic R&D misallocation is less consequential for economic growth because of their dependence on foreign knowledge spillovers.

Table 4 shows the size of the potential welfare gains from adopting the optimal allocation, evaluated using R&D allocation measured in the years 2000, 2005, and 2010. For the year 2010, adopting the optimal allocation in Japan, which has the most efficient R&D allocation in our sample, could lead to welfare improvements equivalent to raising consumption along the entire path

by 5.64%. The potential welfare gains for the U.S. are 8.04% in consumption-equivalent terms,[8] which is above the average in our sample. Russia has the highest potential gains, equivalent to 16.76% consumption gains after adopting the optinal allocation. Moving to the country-specific optimal R&D allocation in 2010 would generate consumption-equivalent welfare gains of 5.60% in China, 4.24% in South Korea, and 4.09% in Germany. For most economies, the size of the potential welfare gains has been stable since the 2000s.

It is important to note that a more allocatively efficient country is not necessarily more innovative in absolute terms. Instead, the extent of misallocation reflects the distance from actual R&D allocation ($\boldsymbol{b}$) to each country's own efficient benchmark ($\boldsymbol{\gamma}$), and our welfare calculations reflect how much each country could gain when moving to that benchmark, holding all other economic conditions fixed.

Table 4. Percentage (%) Consumption Gains By Moving to Each Country's Optimal R&D Allocation

	US	Japan	China	South Korea	Germany	Russia	France	UK	Canada	Netherlands
2000	9.98	4.24	5.78	5.25	4.79	13.70	5.17	7.55	7.22	6.70
2005	8.85	5.04	5.26	3.92	4.11	11.18	5.38	8.17	7.29	5.45
2010	8.04	5.64	5.60	4.24	4.09	16.76	5.38	8.15	6.21	10.22

	Sweden	Switzerland	Italy	Finland	India	Australia	Belgium	Austria	Denmark	European Union
2000	6.65	5.18	5.04	5.39	10.91	5.72	5.72	6.52	5.93	5.91
2005	5.53	4.10	4.57	5.63	8.33	4.19	5.62	8.50	5.30	5.04
2010	6.20	3.67	4.40	7.95	6.21	7.30	6.73	9.87	5.39	5.76

Notes. This table shows the consumption-equivalent welfare gains when each economy moves its optimal R&D allocation, calculated using the formula (31) under our baseline calibration $\{\rho, \lambda, \psi\} = \{0.05, 0.17, 0.06\}$. Section 5.4 discusses robustness under different parameter values.

5.4 Additional Results, Robustness Checks, and Sensitivity Analysis

In Online Appendix E.4, we present additional results and robustness checks on our quantitative analysis of R&D reallocation. Note that our R&D allocation accounting exercise rely on information on the network $\boldsymbol{\Omega}$, the value of the effective discount rate ρ/λ, and the elasticity of productivity to knowledge ψ. In the baseline specification, we compute entry ω_{ij} of $\boldsymbol{\Omega}$ as the share of all citations from i that are towards j ($\frac{Cites_{i\to j}}{\sum_k Cites_{i\to k}}$), and we set $\{\rho, \lambda, \psi\} = \{0.05, 0.17, 0.06\}$. For expositional simplicity, the baseline results focus on top innovative sectors and in the most recent time period. We now tackle the robustness along all these dimensions in the Online Appendix.

[8] In Section B.6 of the Online Appendix, we derive the optimal R&D allocation and the welfare cost formula when the domestic planner takes into account how domestic R&D affects foreign variables. We find very similar welfare gains under that specification.

To start, we reproduce Figure 7 and Table 4 using all IPCs (instead of top 50), and they are shown in Figures A.15 and A.13 and Table A.20. We also present over time allocative efficiency patterns for the 10 countries with the most patents in Figure A.16.

Table A.21 shows that the optimal R&D allocation $\boldsymbol{\gamma}$ is highly stable (by both Pearson's correlation and Spearman's rank correlation) across the following alternative specifications of the innovation network $\boldsymbol{\Omega}$ and parameterizations of ρ and λ.[9]

As discussed in Section 4.1, cross-sector knowledge spillovers are inherently difficult to capture. We consider several alternative specifications of $\boldsymbol{\Omega}$. In rows A1 and A2, we weigh each cite in $\boldsymbol{\Omega}$ construction (24) by the quality (total forward citations received) of either the citing or the cited patent. The innovation networks in these extensions therefore weigh more heavily the spillovers among major patents. In row A3, we construct $\omega_{ij} \propto Cites_{i\to j}$ to scale directly with the total citations totally across or ij-pairs (rather than normalized by the citations from i), and we choose the proportionality constant so that the spectral radius of $\boldsymbol{\Omega}$ is equal to one, ensuring endogenous growth as in our baseline model (see Section B.9 for the theoretical discussion of this specification).

In rows B1 to B7, we consider a range of alternative values for ρ/λ and show that the optimal R&D allocation correlates highly with our baseline specification. We also consider a specification with sector-specific λ_i (row C1). The optimal R&D allocation in this environment is derived in Online Appendix B.8; the heterogeneity in λ_i is measured using variations in each sector's median ROA (return on assets) in our firm-level datasets, with the mapping based on the decentralized economy constructed in Online Appendix B.4.

Finally, one related concern is that citations are only noisy proxies of knowledge spillovers, and thus the innovation network can only be noisy measured. While we cannot purge the measurement error, we conduct the robustness check in reverse and show our quantitative analysis is robust to introducing additional, simulated random errors to $\boldsymbol{\Omega}$, shown in rows D1 to D10.

5.5 How Does R&D Allocation Compare with the Optimal in the U.S.?

We here provide some descriptive evidence for how the actual R&D allocation compares with the optimal in the U.S. Figure 8 plots the log-ratio between the actual R&D expenditure share in the U.S. and the optimal allocation for the 30 largest 3-digit IPC classes by patent output. Altogether, these 30 IPC classes (out of 131) account for 84% of patents and 90% of R&D expenditures in the U.S. Though providing a full set of policy recommendations on R&D allocation is beyond the scope of this paper, this figure conveys several noteworthy messages. Electric communication

[9] We do not report robustness results under alternative values of parameter ψ, as the parameter does not affect the optimal allocation $\boldsymbol{\gamma}$. The welfare impact of R&D allocation is directly proportional ψ, so our baseline results in Table 4 can be directly rescaled for different values of ψ.

technique (H04; e.g., telephonic communication, wireless communication), which ranks 4th in centrality (Figure 2) and 9th in the optimal allocation γ (Figure 4), is over-invested. Meanwhile, within the same broad IPC class H (electricity), the more central and fundamental class Basic Electric Elements (H01; e.g., semiconductor devices) is underinvested. This supports the recent U.S. initiatives (such as the CHIPS For America Act) to invest in the semiconductor industry. Another observation is that the underinvested group (right end of the graph) over-represents IPC classes related to technologies often termed as "green innovation" that can help reduce pollution and the negative consequences of resource exploitation. For instance, the most underinvested IPC class in the figure, B01, covers subclasses on waste management, alternative energy production, and environmental management.[10]

Figure A.17 of the Online Appendix demonstrates the R&D allocation *within* broad technology classes (1-digit IPC) for the U.S. economy. Each of the eight 1-digit IPC categories is represented by a separate panel, in which we show the log-ratio between actual R&D and the constrained-optimal R&D allocation if a planner can reallocate resources across 3-digit IPC classes within the 1-digit IPC category. We find that in IPC category A (i.e., human necessities), medical or veterinary science (A61) is over-allocated with R&D relative to agriculture (A01) and foods (A23); in IPC category C (chemistry; metallurgy), organic chemistry (C07) is over-allocated relative to petroleum, gas, or coke industries (C10); in category E (fixed construction), locks, keys, and safes (E05) is over-allocated relative to mining (E21).

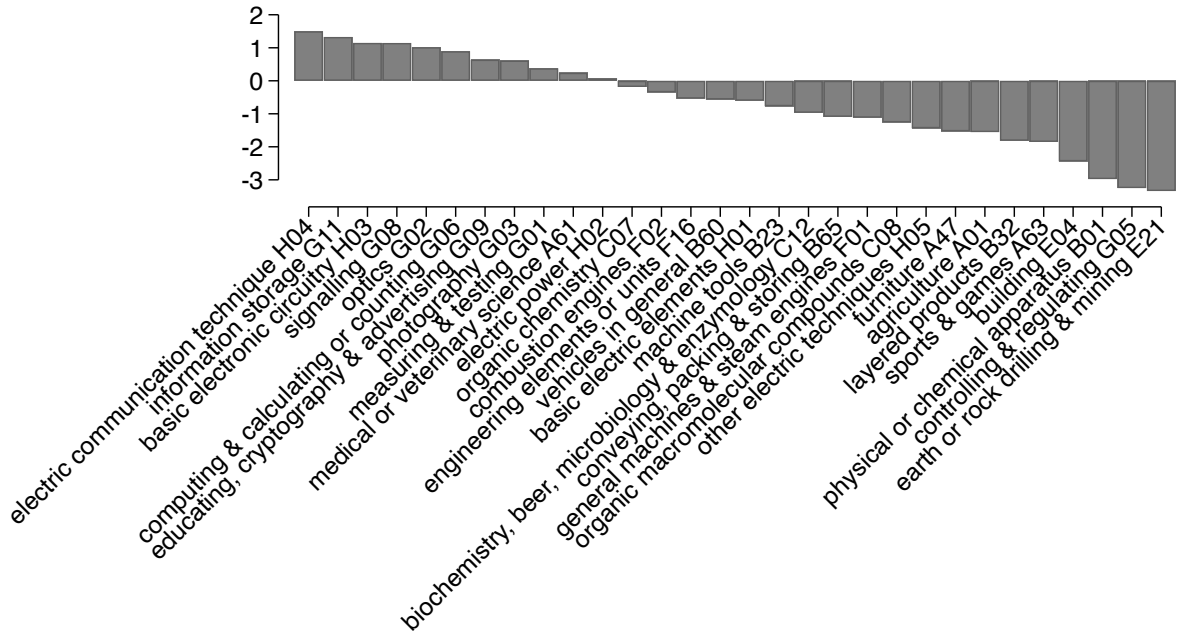
5.6 Innovation Hubs

What explains cross-country differences in R&D allocative efficiency? We do not have definitive answers, but we can present a conjecture with some empirical support: firms whose R&D activities span multiple sectors and technology classes allocate their resources in ways that may resemble the social planner's. Because these firms' R&D activities build on their own prior innovations, they may partially internalize knowledge spillovers through the innovation network. Notable examples include top innovating firms such as IBM, Samsung, Sony, and Siemens, which are termed "innovation hubs."

Our hypothesis is supported empirically by a strong negative relationship between the presence of such firms and the degree of R&D misallocation in each country. Figure 9, Panel (a) shows the share of patents in 2010–2014 that are filed by the top 10% of innovating firms in each country. The figure shows that R&D activities are more concentrated in Japan, the U.S., and Sweden, as the top 10% of innovating firms in these economies account for close to 90%, 80%, and 70% of patents,

[10]For more references on identifying green innovation-related IPC classes, see Cohen, Gurun, and Nguyen (2020).

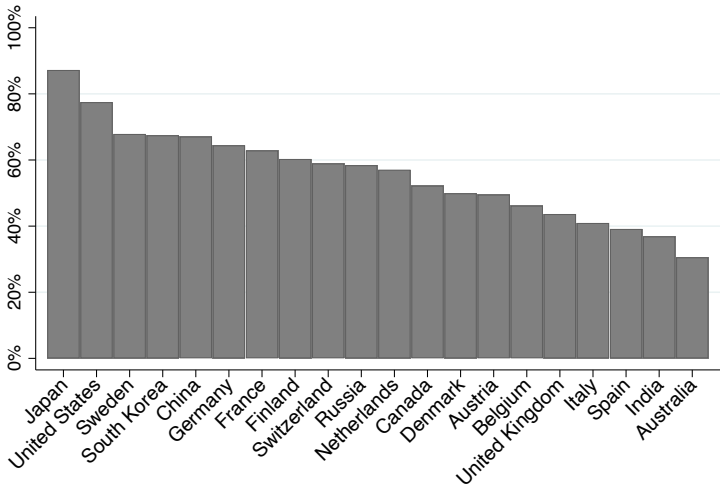
Figure 8. U.S. R&D Misallocation in the Top 30 Innovative IPC Classes



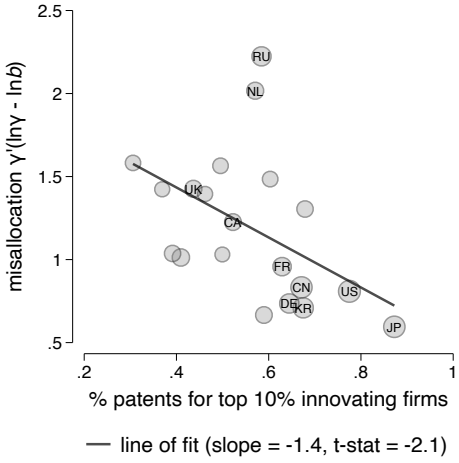
Notes. This figure plots the level of misallocation of the top 30 innovative IPC classes, ranked using total patent output. The level of misallocation is calculated as $\ln \boldsymbol{b} - \ln \boldsymbol{\gamma}$. Positive bars (left end) mean over-investment, and negative bars mean underinvestment.

Figure 9. Innovation Hubs and R&D Allocative Efficiency

(a) Share of Patents from Top Innovative Firms



(b) Innovation Hubs and R&D Allocative Efficiency



Notes. Panel (a) of this figure shows the share of patents filed by the top 10% of innovative firms in each country between 2010–2014 (innovative firms are ranked using patent output). Panel (b) plots the misallocation measure against the measure of concentration in Panel (a).

respectively. By contrast, R&D activities are least concentrated in Spain, India, and Australia.

Panel (b) of Figure 9 plots the misallocation measure ($\boldsymbol{\gamma}' (\ln \boldsymbol{\gamma} - \ln \boldsymbol{b})$) against the share of

patents accounted for by the top 10% of innovating firms. We find a strongly negative relationship (slope -1.4, t-statistic -2.1). This evidence suggests that the market failure in R&D resource allocation could be partially mitigated if innovation hub firms thrive.

6 Conclusion

We study optimal cross-sector allocation of R&D resources in an endogenous growth model featuring an innovation network. We provide closed-form solutions for the optimal path of R&D resource allocation, and we show a planner valuing long-term growth (i.e., with low discount rates) should allocate more R&D toward key sectors that are central in the innovation network, but the incentive is muted in open economies that benefit more from foreign knowledge spillovers. We show the relative entropy of actual R&D allocation from the optimal allocation maps into a sufficient statistic for the potential welfare gains from reallocating R&D optimally.

To empirically evaluate R&D allocative efficiency across countries and over time, we build a global innovation network based on over 30 million global patents and compile comprehensive data on sectoral production, final use, and, importantly, R&D resource allocation for major innovative economies. We find that our model-implied optimal R&D resource allocation explains real allocations in the data, particularly for countries generally perceived as innovative, such as the U.S., Japan, Germany, and more recently China and South Korea. However, there remains significant room to improve. Improving R&D allocations could generate substantial welfare improvements across the globe. For the U.S., reallocating R&D resources to Japan's efficiency level would increase consumption-equivalent welfare by 19.6% in 2010. We believe our framework can be adopted to explore future questions about R&D allocation.

References

Acemoglu, D., U. Akcigit, N. Bloom, and W. R. Kerr (2018): "Innovation, reallocation and growth," *American Economic Review*, 108, 3450–3491.

Acemoglu, D., U. Akcigit, and W. Kerr (2015): "Networks and the Macroeconomy: An Empirical Exploration," 273–335.

Acemoglu, D., U. Akcigit, and W. R. Kerr (2016): "Innovation network," *Proceedings of the National Academy of Sciences*, 113, 11483–11488.

Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi (2012): "The Network Origins of Aggregate Fluctuations," *Econometrica*, 80, 1977–2016.

Aghion, P., N. Bloom, R. Blundell, R. Griffith, and P. Howitt (2005): "Competition and innovation: An inverted-U relationship," *Quarterly Journal of Economics*, 120, 701–728.

Akcigit, U., S. Ates, and G. Impullitti (2022): "Innovation and Trade Policy in a Globalized World," *Working paper*.

Akcigit, U., D. Hanley, and N. Serrano-Velarde (2021): "Back to basics: Basic research spillovers, innovation policy, and growth," *Review of Economic Studies*, 88, 1–43.

Akcigit, U. and W. R. Kerr (2018): "Growth through heterogeneous innovations," *Journal of Political Economy*, 126, 1374–1443.

Atkeson, A. and A. Burstein (2019): "Aggregate implications of innovation policy," *Journal of Political Economy*, 127, 2625–2683.

Baqaee, D. R. (2018): "Cascading Failures in Production Networks," *Econometrica*, 86.

Baqaee, D. R. and E. Farhi (2019): "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem," *Econometrica*, 87.

——— (2020): "Productivity and Misallocation in General Equilibrium," *The Quarterly Journal of Economics*, 135, 105–163.

Bloom, N., C. I. Jones, J. Van Reenen, and M. Webb (2020): "Are ideas getting harder to find?" *American Economic Review*, 110, 1104–44.

Bloom, N., M. Schankerman, and J. Van Reenen (2013): "Identifying technology spillovers and product market rivalry," *Econometrica*, 81, 1347–1393.

Bloom, N., J. Van Reenen, and H. Williams (2019): "A toolkit of policies to promote innovation," *Journal of Economic Perspectives*, 33, 163–84.

Buera, P. and E. Oberfield (2020): "The Global Diffusion of Ideas," *Econometrica*.

Caballero, R. J. and A. B. Jaffe (1993): "How High Are the Giants' Shoulders: An Empirical Assessment of Knowledge Spillovers and Creative Destruction in a Model of Economic Growth," *NBER Macroeconomics Annual*, 8, 15–74.

Cai, J. and N. Li (2019): "Growth Through Intersectoral Knowledge Linkages," *Review of Economic Studies*.

Cai, J., N. Li, and A. M. Santacreu (2022): "Knowledge Diffusion, Trade and Innovation across Countries and Sectors," *AEJ: Macroeconomics*, 14, 104–45.

Cai, J. and C. Tian (2021): "Innovation and Endogenous Knowledge Network Dynamics," *Working paper*.

Carvalho, V. (2010): "Aggregate fluctuations and the network structure of intersectoral trade," *Working paper*.

Chaney, T. (2018): "The Gravity Equation in International Trade: An Explanation," *Journal of Political Economy*.

Coe, D. T. and E. Helpman (1995): "International R&D spillovers," *European Economic Review*.

Coe, D. T., E. Helpman, and A. W. Hoffmaister (2009): "International R&D Spillovers and Institutions," *European Economic Review*.

Cohen, L., U. G. Gurun, and Q. H. Nguyen (2020): "The ESG-innovation disconnect: Evidence from green patenting," .

David, J. and V. Venkateswaran (2019): "The Sources of Capital Misallocation," *American Economic Review*.

Eaton, J. and S. Kortum (1999): "International Technology Diffusion: Theory and Measurement," *International Economic Review*.

——— (2006): "Innovation, Diffusion, and Trade," *NBER Working paper*.

Gabaix, X. (2011): "The Granular Origins of Aggregate Fluctuations," *Econometrica*.

Galeotti, A., B. Golub, and S. Goyal (2020): "Targeting Interventions in Networks," *Econometrica*.

Garcia-Macia, D., C.-T. Hsieh, and P. J. Klenow (2019): "How Destructive Is Innovation?" *Econometrica*, 87, 1507–1541.

Grassi, B. (2017): "IO in I-O: Size, Industrial Organization, and the Input-Output Network Make a Firm Structurally Important," *Working Paper*.

Griffith, R., R. Harrison, and J. Van Reenen (2006): "How special is the special relationship? Using the impact of US R&D spillovers on UK firms as a test of technology sourcing," *American Economic Review*, 96, 1859–1875.

Guillard, C., R. Martin, P. Mohnen, C. Thomas, and V. Dennis (2021): "Efficient industrial policy for innovation: standing on the shoulders of hidden giants," *Working paper*.

Hopenhayn, H. and F. Squintani (2021): "On the Direction of Innovation," *Journal of Political Economy*.

Hsieh, C.-T., E. Hurst, C. Jones, and P. Klenow (2019): "The Allocation of Talent and U.S. Economic Growth," *Econometrica*.

Hsieh, C.-T. and P. Klenow (2009): "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*.

Huang, J. and Y. Zenou (2020): "Key Sectors in Endogeneous Growth," *Working paper*.

Hulten, C. R. (1978): "Growth Accounting with Intermediate Inputs," *The Review of Economic Studies*, 45, 511–518.

Jaffe, A. B., M. Trajtenberg, and R. Henderson (1993): "Geographic localization of knowledge spillovers as evidenced by patent citations," *Quarterly Journal of Economics*, 108, 577–598.

Jones, B. F. and L. H. Summers (forthcoming): "A Calculation of the Social Returns to Innovation," in *Innovation and Public Policy*, ed. by B. F. Jones and A. Goolsbee, NBER.

Jones, C. (2022): "The Past and Future of Economic Growth: A Semi-Endogenous Perspective," *Annual Review of Economics*.

Jones, C. I. (1995): "R&D-Based Models of Economic Growth," *Journal of Political Economy*.

——— (2011): "Intermediate Goods and Weak Links in the Theory of Economic Development," *American Economic Journal: Macroeconomics*, 3, 1–28.

——— (2013): "Misallocation, Economic Growth, and Input-Output Economics," *Advances in Economics and Econometrics*, 2.

Jones, C. I. and J. Williams (1998): "Measuring the Social Return to R&D," *Quarterly Journal of Economics*.

Keller, W. (2004): "International Technology Diffusion," *Journal of Economic Literature*.

Kleinman, B., E. Liu, and S. Redding (2022): "International Friends and Enemies," *Working Paper*.

——— (forthcoming): "Dynamic Spatial General Equilibrium," *Econometrica*.

Koenig, M., M. Z. Song, K. Storesletten, and F. Zilibotti (forthcoming): "From Imitation to Innovation: Where Is All that Chinese R&D Going?" *Econometrica*.

Kogan, L., D. Papanikolaou, A. Seru, and N. Stoffman (2017): "Technological innovation, resource allocation, and growth," *Quarterly Journal of Economics*, 132, 665–712.

Lim, K. (2018): "Endogenous Production Networks and the Business Cycle," *Working Paper*.

Liu, E. (2019): "Industrial Policies in Production Networks*," *The Quarterly Journal of Economics*, 134, 1883–1948.

Liu, E. and A. Tsyvinski (2022): "Dynamical Structures and Spectral Properties of Input-Output Networks," *Working Paper*.

Lucking, B., N. Bloom, and J. Van Reenen (2018): "Have R&D Spillovers Changed?" *NBER Working Paper*.

Lybbert, T. J. and N. J. Zolas (2014): "Getting patents and economic data to speak to each other: An algorithmic links with probabilities approach for joint analyses of patenting and economic activity," *Research Policy*, 43, 530–542.

Ma, S. (2020): "The life cycle of corporate venture capital," *Review of Financial Studies*, 33, 358–394.

——— (2021): "Technological Obsolescence," *Working Paper*.

Machin, S. and J. Van Reenen (1998): "Technology and changes in skill structure: evidence from seven OECD countries," *The quarterly journal of economics*, 113, 1215–1244.

Manski, C. F. (1993): "Identification of endogenous social effects: The reflection problem," *Review of Economic Studies*, 60, 531–542.

Melitz, M. and S. Redding (2021): "Trade and Innovation," *The Economics of Creative Destruction, Festschrift Volume in honor of Philippe Aghion and Peter Howitt*.

Oberfield, E. (2018): "A Theory of Input-Output Architecture," *Econometrica*, 86.

Restuccia, D. and R. Rogerson (2008): "Policy distortions and aggregate productivity with heterogeneous establishment," *Review of Economic Dynamics*.

Santacreu, A. M. (2015): "Innovation, Diffusion, and Trade: Theory and Measurement," *Journal of Monetary Economics*.

Taschereau-Dumouchel, M. (2020): "Cascades and Fluctuations in an Economy with an Endogenous Production Network," *Working Paper*.

Thomson, R. (2017): "The effectiveness of R&D tax credits," *Review of Economics and Statistics*, 99, 544–549.

Timmer, M. P., E. Dietzenbacher, B. Los, R. Stehrer, and G. J. de Vries (2015): "An Illustrated User Guide to the World Input-Output Database: the Case of Global Automotive Production," *Review of International Economics*.

vom Lehn, C. and T. Winberry (2022): "The Investment Network, Sectoral Comovement, and the Changing U.S. Business Cycle," *Quarterly Journal of Economics*.

Wilson, D. J. (2009): "Beggar thy neighbor? The in-state, out-of-state, and aggregate effects of R&D tax credits," *The Review of Economics and Statistics*, 91, 431–436.